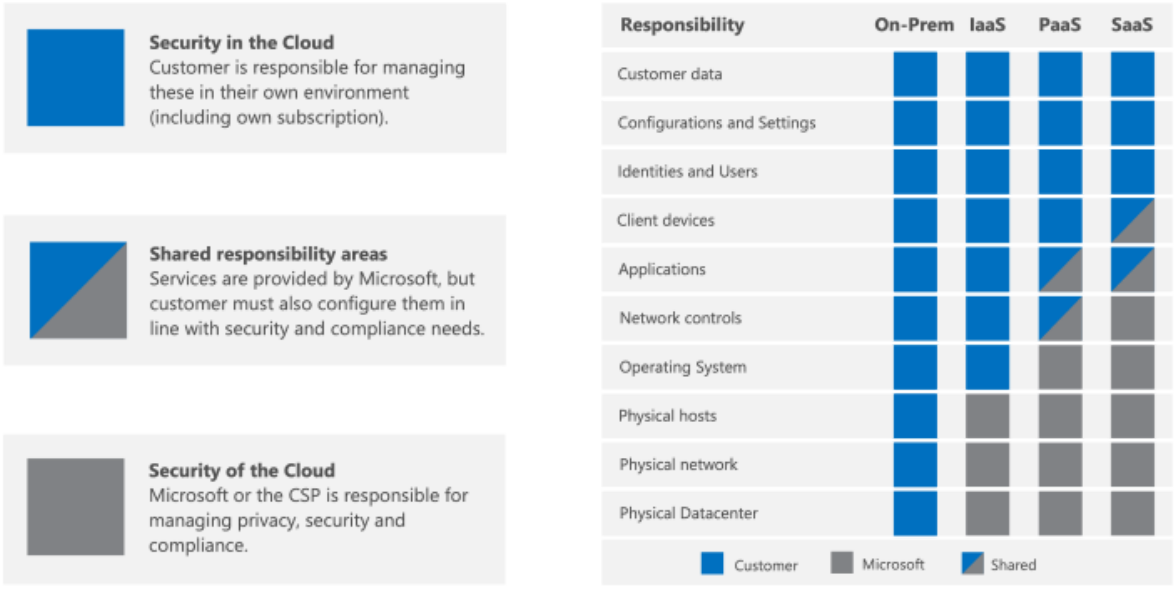
Tell us about your PDF experience.

Risk Assessment Guide for Microsoft Cloud

Article • 03/02/2023

The goal of a cloud risk assessment is to ensure that the system and data considered for migration to the cloud don't introduce any new or unidentified risk into the organization. The focus is to ensure confidentiality, integrity, availability, and privacy of information processing and to keep identified risks below the accepted internal risk threshold.

In a shared responsibility model, the Cloud Service Provider (CSP) is responsible for managing security and compliance *of the cloud* as the provider. The customer remains responsible for managing and configuring security and compliance *in the cloud* in accordance with their needs and risk tolerance.



In this guide, best practices are shared on how to efficiently evaluate vendor risks and how to use the resources and tools Microsoft makes available.

Understand shared responsibility in the cloud

Cloud deployments can be categorized as Infrastructure as a Service (IaaS), Platform as a Service (PaaS), or Software as a Service (SaaS). Depending upon the applicable cloud service model, the level of responsibility over the solutions' security controls shift between the CSP and the customer. In a traditional on-premises model, the customer is responsible for the whole stack. When moving to the cloud, all physical security

responsibilities transfer to the CSP. Depending on the cloud service model for your organization, additional responsibilities shift over to the CSP. However, in most service models, your organization remains responsible for the devices used to access the cloud, network connectivity, your accounts and identities, and your data. Microsoft invests heavily in creating services that allow customers to stay in control of their data across the entire lifecycle.

Microsoft Cloud operates at a hyperscale, relying on a combination of DevSecOps and automation to standardize operating models. Microsoft operating model changes the way risk is approached compared to traditional on-premises operating models, leading to the implementation of different and sometimes unfamiliar controls to manage risks. When conducting your cloud risk assessment, keep in mind that Microsoft's goal is to ensure all risks are addressed, but not necessarily to implement the same controls your organization does. Microsoft may address the same risks with a different set of controls and that should be reflected in the cloud risk assessment. Designing and implementing strong preventive controls can reduce much of the work required by the detective and corrective controls. An example of this is Microsoft's implementation of Zero Standing Access (ZSA).

Adopt a framework

Microsoft recommends that customers map their internal risk and controls framework to an independent framework that addresses cloud risks in a standardized way. If your existing internal risk assessment models don't address the specific challenges that come with cloud computing, you'll benefit from these broadly adopted and standardized frameworks. A secondary benefit is that Microsoft provides mappings against these frameworks in documentation and tools that will accelerate your risk assessments. Examples of these frameworks include the ISO 27001 Information security standard, CIS Benchmark, and NIST SP 800-53. Microsoft offers the most comprehensive set of compliance offerings of any CSP. For more information, see Microsoft compliance offerings.

Use Microsoft Purview Compliance Manager to create your own assessments that evaluate compliance with the industry and regional regulations that apply to your organization. Assessments are built upon the framework of assessment templates, which contain the necessary controls, improvement actions, and where applicable, Microsoft actions for completing the assessment. For Microsoft actions, detailed implementation plans and recent audit results are provided. This way, time can be saved on fact finding, mapping, and researching how specific controls are implemented by Microsoft. For more information, see the Microsoft Purview Compliance Manager article.

Understand how Microsoft operates to safeguard your data

While the customer is responsible for managing and configuring security and compliance *in the cloud*, the CSP is responsible for managing security and compliance *of the cloud*. One way to validate that the CSP is effectively addressing their responsibilities and upholding their promises is to review their external audit reports such as ISO and SOC. Microsoft makes external audit reports available to authenticated audiences on the Service Trust Portal (STP).

In addition to external audit reports, Microsoft highly encourages customers to take advantage of the following resources to help understand how Microsoft operates in depth:

- On-demand learning path: Microsoft Learn offers hundreds of learning paths and modules on different topics. Amongst them, take Learn how Microsoft safeguards customer data to understand Microsoft's fundamental security and privacy practices.
- Service Assurance on Microsoft Compliance: Articles on Microsoft's practices are categorized into 16 domains for easier review. Each domain includes an overview that captures how Microsoft is managing risks associated with each area. Audit tables are provided containing links to the most recent reports stored on the STP, related sections, and the date the audit report was conducted for Microsoft online services. If available, links to artifacts demonstrating control implementation, such as third-party vulnerability assessments and business continuity plan verification reports are provided. Like audit reports, these artifacts are hosted on STP and require authentication to access.

Domain	Description
Architecture	The design of Microsoft online services and the security principles that act as its foundation.
Audit logging and monitoring	How Microsoft captures, processes, stores, protects, and analyzes logs to detect unauthorized activity and monitor performance. make security and performance monitoring possible.
Datacenter security	How Microsoft securely operates the datacenters that provide the means to operate Microsoft online services worldwide.
Encryption and key management	The cryptographic protection of customer communications and the data stored and processed in the cloud.

Domain	Description
Governance, risk, and compliance	How Microsoft enforces the security policies it creates and manages risk to meet customer promises and compliance requirements.
Identity and access management	The protection of Microsoft online services and customer data from unauthorized or malicious access.
Security incident management	The processes Microsoft uses to prepare for, detect, respond, and communicate all security incidents.
Network security	How Microsoft protects its network boundaries from external attacks and manages its internal network to limit their propagation.
Personnel management	The screening processes, training, and secure management of personnel throughout their time at Microsoft.
Privacy and data management	How Microsoft handles and protects customer data to preserve their data rights.
Resiliency and continuity	Process and technologies used to maintain service availability and ensure business continuity and recovery.
Security development and operation	How Microsoft ensures that its services are designed, run, and managed securely throughout their lifecycle.
Supplier management	How Microsoft screens and manages third-party companies that assist with Microsoft online services.
Threat and vulnerability management	The processes Microsoft uses to scan for, detect, and address vulnerabilities and malware.

Compliance Program for Microsoft Cloud (CPMC)

Organizations share responsibility with their CSP to protect the data and systems that exist in the cloud. Customers need to assess the risks and address regulatory compliance requirements in an efficient, repeatable manner. However, it can be difficult to navigate the global regulatory landscape and gain enough insight into a CSP's practices to reach an acceptable level of assurance. To overcome these challenges, Microsoft launched the Compliance Program for Microsoft Cloud (CPMC). The CPMC is a fee-based premium

program offering personalized regulatory and industry specific compliance support, education, and networking opportunities. For more information on CPMC specific offerings, check out the CPMC website.

Resources

- Video: Learn how Microsoft safeguards customer data
- Cloud Governance utilizing Microsoft Hyper Scale Cloud Computing (HSCC) solutions
- Risk Assessment and Compliance Guide for Financial Institutions in the Microsoft Cloud
- Concentration Risk: Perspectives from Microsoft
- Service Trust Portal

Architecture overview

Article • 04/04/2023

What are Microsoft online services?

Microsoft online services refer to the cloud-based services offered by Microsoft, which includes Azure, Dynamics 365, and Microsoft 365. Each service offers unique solutions to common business operation and productivity needs, serving customers around the world ranging from small businesses to large enterprises and governments. Globally distributed datacenters connected by Microsoft's independently managed network infrastructure act as the backbone to support online services, providing the ability to sustain availability in almost every situation and scale to massive worldwide demand.

All Microsoft online services have the same goal of safeguarding the service infrastructure they manage and their customers' data, but each online service is operated by a separate business unit. In many instances, security controls are implemented in the same way across all services, but in some cases, they take a different approach to fulfilling their customer promises and compliance obligations.

What is Azure?

Microsoft Azure is a cloud computing platform for building, deploying, and managing applications through a global network of Microsoft and third-party managed datacenters. It supports both Platform as a Service (PaaS) and Infrastructure as a Service (IaaS) cloud service models, and enables hybrid solutions that integrate cloud services with customers' on-premises resources. Microsoft Azure supports many customers, partners, and government organizations that span across a broad range of products and services, geographies, and industries. Microsoft Azure is designed to meet their security, confidentiality, and compliance requirements.

What is Dynamics 365?

Dynamics 365 is an online business application suite that integrates the Customer Relationship Management (CRM) capabilities and its extensions with the Enterprise Resource Planning (ERP) capabilities. These end-to-end business applications help customers turn relationships into revenue, earn customers, and accelerate business growth. Dynamics 365 is a Software as a Service (SaaS) suite built on Azure's

infrastructure and is made available to customers around the world through their globally distributed datacenters.

What is Microsoft 365?

Microsoft 365 is the cloud-powered, subscription-based version of Office, Windows 10, Enterprise Mobility + Security, and Compliance. Microsoft 365 customers get clients such as Outlook and Windows, and they also benefit from services that Microsoft hosts on their behalf, such as Exchange Online, Microsoft Teams, and SharePoint Online. All components of the service are regularly updated as part of the subscription model, so that our customers have an 'evergreen' product. Microsoft manages the service infrastructure on behalf of customers, meaning that Microsoft is responsible for securing the infrastructure that stores customer data.

In terms of scale, Microsoft currently uses close to a million machines to power Microsoft 365 services. The infrastructure powering these services varies widely across service-specific hardware and virtualized environments in Azure, Windows and Linux, and multi-tenant and dedicated platforms. Microsoft 365 is a global business, and our infrastructure is distributed in datacenters around the world, enabling our customers to meet data residency and sovereignty requirements.

How do Microsoft online services ensure isolation between customer tenants?

Microsoft's cloud services are built on the assumption that all tenants are potentially hostile to all other tenants. To properly isolate tenants from one another, Microsoft implements various isolation technologies and controls. These controls are designed to safeguard against information leakage or unauthorized access to customer data across tenants and to prevent the actions of one tenant from adversely affecting the service for another tenant.

Customer content is logically isolated within tenants using Microsoft Entra ID. User authentication in Microsoft online services verifies not only the user identity, but also the tenant identity the user account is part of, preventing users from accessing data outside their tenant environment. To supplement the logical isolation of Microsoft Entra ID, customer content is always encrypted at rest and in transit. Individual services may also provide additional layers of tenant isolation, such as SharePoint Online isolation of tenant data in separate, encrypted databases.

How do Microsoft online services engineer resilient services that avoid single points of failure?

Microsoft designs and builds cloud services to maximize reliability and minimize the negative effects on customers in the face of faults and challenges to normal operations. This strategy begins with the design of the network connecting our geographically distributed datacenters. Microsoft's network architecture includes direct interconnections and multiple network paths. Microsoft online services use this redundancy to automatically route traffic around failures to improve service quality.

Wherever possible, Microsoft online services are deployed in active/active configurations with automated service health monitoring, allowing the service to detect and recover from many common faults and failures without human intervention. In addition to active/active configurations, Microsoft online services increase fault-tolerance by ensuring the service is deployed in separate fault zones, preventing a fault in one zone from affecting the availability of other zones.

Data resiliency complements service resiliency by protecting the integrity and availability of data in Microsoft online services. Our services use local storage redundancy and geo-redundancy to replicate copies of customer data into different fault zones. If data is corrupted or lost in one fault zone, it can be accessed in another fault zone without loss of availability. Automated integrity checking automatically restores data impacted by many kinds of physical or logical corruption. Microsoft also provides customers with tools to restore data accidentally deleted or modified by the customer in services such as Exchange Online and SharePoint Online.

How do Microsoft online services track dependencies and prevent unauthorized external system connections?

Microsoft online services teams identify critical system components and their dependencies as part of Business Continuity Management. In addition, Microsoft documents and tracks all external system connections to ensure only authorized connections are allowed in network firewall configurations. Microsoft online services systems, dependencies, and external connections are documented in Microsoft online services' information security architecture. Both the information security architecture and corresponding data flow diagrams are reviewed and updated annually at a minimum, as well as whenever significant changes are made to the system.

Microsoft online services architecture is validated regularly and automatically using cloud-based tools to verify alignment with our security principles and to continuously test isolation and resiliency features. Architectural validation works to automatically identify instances where the current state of the service has drifted from the desired state, flagging any deviations for review and mitigation. The goal of architecture validation is to ensure the security capabilities of our service infrastructure continue to function as expected.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to architecture.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
ISO 27001/27002 Statement of Applicability Certificate	A.6: Organization of information security A.13.1: Network security management A.17.2: Redundancies	April 24, 2023
ISO 27017 Statement of Applicability Certificate	A.6: Organization of information security A.13.1: Network security management A.17.2: Redundancies	April 24, 2023
SOC 1 SOC 2 SOC 3	BC-1: Business continuity plans BC-3: BCDR procedures BC-4: BCDR testing BC-5: Business continuity risk assessments BC-6: Third-party business continuity BC-7: Datacenter business continuity procedures BC-8: Datacenter business continuity testing BC-9: Datacenter resiliency assessment DS-6: Redundancy of critical components DS-7: Replication of customer data DS-14: Restoration of customer services DS-16: Customer data segregation	November 17, 2023

Microsoft 365

Expand table

External audits	Section	Latest report date
FedRAMP (Office 365)	AC-4: Information flow enforcement CP-9: Information system backup PL-8: Information security architecture SC-7: Boundary protection SC-22: Architecture and provisioning	July 31, 2023
ISO 27001/27002/27017 Statement of Applicability Certification (27001/27002) Certification (27017)	A.6: Organization of information security A.13.1: Network security management A.17.2: Redundancies	March 2023
SOC 1	CA-37: Tenant isolation CA-49: Backup policies CA-51: Data replication	January 3, 2023
SOC 2	CA-05: Data flow diagrams CA-37: Tenant isolation CA-49: Backup policies CA-51: Data replication	January 3, 2023

Resources

- Interactive: Global Azure Infrastructure

Feedback

Was this page helpful? Yes No

Securing the Microsoft Online Services infrastructure

Article • 06/05/2023

Microsoft Online Services are some of the largest enterprise and consumer cloud services in the world and continue to grow rapidly in their customer base, products, and features. Customers turn to Microsoft not only for its world-class cloud offerings, but to help protect their most sensitive information from the constantly evolving cyber threat landscape. It is Microsoft's top priority to keep customer data secure and maintain customer trust.

Securing a system of this scale and complexity is not possible if security is an afterthought, it is only effective if security is integrated during the initial design process. It requires a robust threat detection system with prompt responses from both automated systems and highly skilled engineers. Continuous assessment and validation of these systems is essential to ensure secure configurations remain intact and previously unknown vulnerabilities are identified.

Core security principles

Seven security principles lay the foundation for our framework of *protecting* Microsoft Online Services from threats, *detecting and responding* to any threats, continuously *assessing* the security posture and improving services based on the results of those assessments.

- **Data privacy**: Customers own their data and Microsoft is the custodian. Microsoft Online Services are designed to operate without engineers accessing customer data, unless explicitly requested by the customer.
- **Assume breach**: Personnel and services are treated as though compromise is a real possibility.
- **Least privilege**: Access and permissions to resources are limited to only what is necessary to perform needed tasks.
- **Breach boundaries**: Identities and infrastructure in one boundary are isolated from resources in other boundaries. Compromise of one boundary should not lead to compromise of another.
- **Service fabric integrated security**: Security priorities and requirements are built into the design of new features and capabilities, ensuring that a strong security posture scales with each service.

- **Automated and automatic**: Microsoft focuses on developing durable products and architectures that can intelligently and automatically enforce service security, while giving Microsoft engineers the ability to safely manage responses to security threats at scale.
- **Adaptive security**: Microsoft security capabilities adapt to and are enhanced by machine learning models, routine penetration testing, and automated assessments.

Protection

Access control

By default, personnel responsible for developing and maintaining Microsoft Online Services have Zero Standing Access (ZSA) to the service infrastructure. While Microsoft strives to hire only the best engineers and rigorous background checks are required, Microsoft does not assume that they are trusted by default in operating services. Additionally, when engineers are approved for privileged access, they are only granted access for a limited duration to perform only the actions needed for a specific scope of the service infrastructure. Microsoft refers to these policies as Just-in-Time (JIT) and Just-Enough-Access (JEA) which are implemented through a system called Lockbox.

To acquire elevated privileges, Microsoft engineers submit a request for the specific task and specify the time frame to perform it. Once approved, Lockbox generates a specialized JIT account with the ability to perform only the requested task. Actions usually take the form of automated workflows that securely perform any troubleshooting or recovery required. In rare instances when direct access to the infrastructure is necessary, strictly monitored Secure Admin Workstations (SAWs) are required.

Rogue users and compromised accounts are a real possibility in any organization, and our access control system are designed to protect against these threats.

For more information about access control, see Identity and access management overview.

Encryption

While access controls provide a vital role in defending Microsoft Online Services services, encryption is used throughout the data lifecycle to further protect confidentiality and privacy for Microsoft customers.

Data in transit between client machines, Microsoft Online Services servers, and non-Microsoft servers is encrypted using TLS 1.2. We regularly review the ciphers and protocols in use, adding improved protocols when available and removing weaker ones as needed.

Customer content at rest on Microsoft servers is encrypted at the volume-level using BitLocker. Application-level encryption can additionally be applied using keys managed by either Microsoft or the customer. Access to Microsoft-managed keys is only possible when authorized and approved through the JIT and JEA process.

For more information about encryption, see Encryption and key management overview.

Network isolation

In line with the principle of least privilege, Microsoft restricts communication between different parts of the service infrastructure to only what is necessary to operate. All network traffic is denied by default, with only explicitly defined communication being allowed. This restriction establishes breach boundaries throughout the infrastructure. Teams that would like to add new network paths to accommodate a new feature to their service must have the request evaluated and approved before it can be opened.

For more information about network isolation, see Microsoft 365 isolation controls and Azure isolation controls.

Detection & Response

Security monitoring

Security monitoring at Microsoft's massive scale is only possible by generating highly accurate alerts using automated cloud-based solutions. Audit logs from each service and telemetry data gathered from throughout the core infrastructure is sent to a proprietary centralized near-real-time processing and alerting solution.

Detected threats are remediated using automatically triggered actions when possible. When automated solutions are unsuccessful or incapable of resolving the issue, on-call Microsoft engineers immediately take action to mitigate the threat.

For more information about security monitoring, see Security monitoring overview.

Assessment

Automated assessments

Regardless of how a system is designed, the security posture can degrade due to intentional and unintentional configuration drift over time. Automated tools constantly assess Microsoft Online Services systems that look for unpatched and misconfigured services. This assessment is often referred to as patching, anti-virus, vulnerability, and configuration scanning (PAVC).

Our architecture is also frequently validated, identifying instances such as unused open ports and accounts with standing administrative access. Any services that drift from a pre-defined desired state are automatically brought back into alignment.

For more information about vulnerability management, see Vulnerability management overview.

Attack simulation and penetration testing

Microsoft's top priority is to prevent attacks from infiltrating defenses. Microsoft Online Services have dedicated teams of security experts who are constantly conducting simulated attacks to identify previously unknown vulnerabilities and to provide a constant stream of rich data to improve security monitoring capabilities. These simulated attacks take the form of frequent automated small-scale attacks and expert-driven deep dives. From these activities, Microsoft evaluates the ability to detect, respond, and evict attackers.

For more information about penetration testing, see Attack simulation in Microsoft 365.

Resources

Behind the Scenes: Securing the Infrastructure Powering the Microsoft 365 Service

Isolation in the Azure Public Cloud

Article • 10/12/2023

Azure allows you to run applications and virtual machines (VMs) on shared physical infrastructure. One of the prime economic motivations to running applications in a cloud environment is the ability to distribute the cost of shared resources among multiple customers. This practice of multi-tenancy improves efficiency by multiplexing resources among disparate customers at low costs. Unfortunately, it also introduces the risk of sharing physical servers and other infrastructure resources to run your sensitive applications and VMs that may belong to an arbitrary and potentially malicious user.

This article outlines how Azure provides isolation against both malicious and non-malicious users and serves as a guide for architecting cloud solutions by offering various isolation choices to architects.

Tenant Level Isolation

One of the primary benefits of cloud computing is concept of a shared, common infrastructure across numerous customers simultaneously, leading to economies of scale. This concept is called multi-tenancy. Microsoft works continuously to ensure that the multi-tenant architecture of Microsoft Cloud Azure supports security, confidentiality, privacy, integrity, and availability standards.

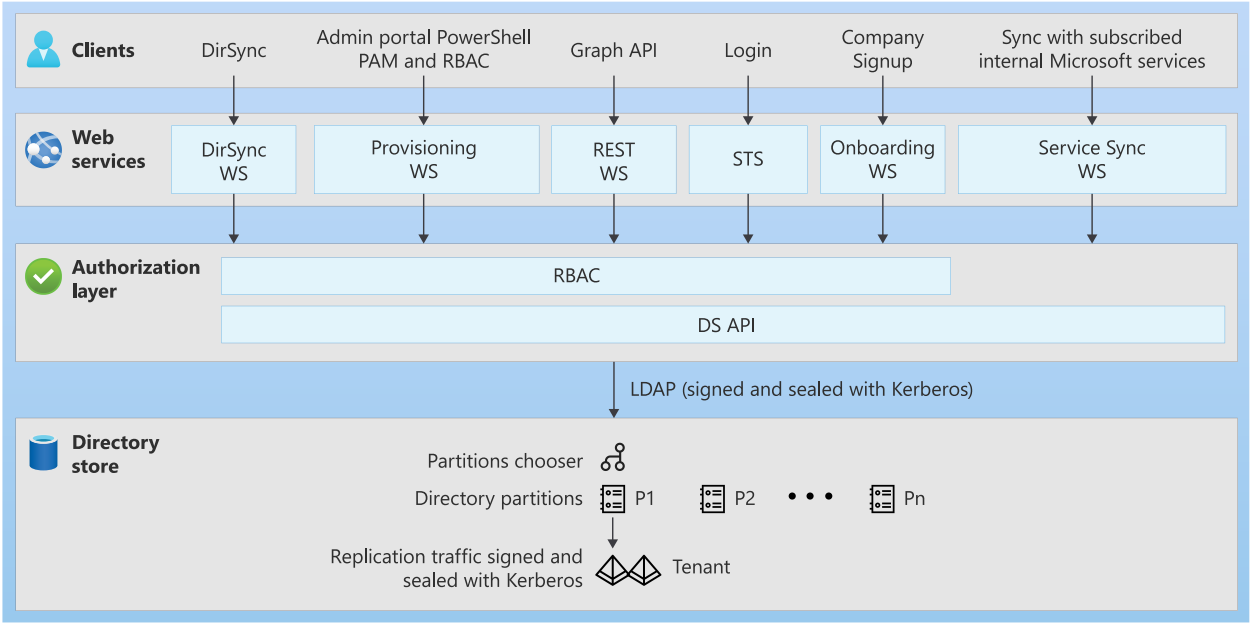
In the cloud-enabled workplace, a tenant can be defined as a client or organization that owns and manages a specific instance of that cloud service. With the identity platform provided by Microsoft Azure, a tenant is simply a dedicated instance of Microsoft Entra ID that your organization receives and owns when it signs up for a Microsoft cloud service.

Each Microsoft Entra directory is distinct and separate from other Microsoft Entra directories. Just like a corporate office building is a secure asset specific to only your organization, a Microsoft Entra directory was also designed to be a secure asset for use by only your organization. The Microsoft Entra architecture isolates customer data and identity information from co-mingling. This means that users and administrators of one Microsoft Entra directory can't accidentally or maliciously access data in another directory.

Azure Tenancy

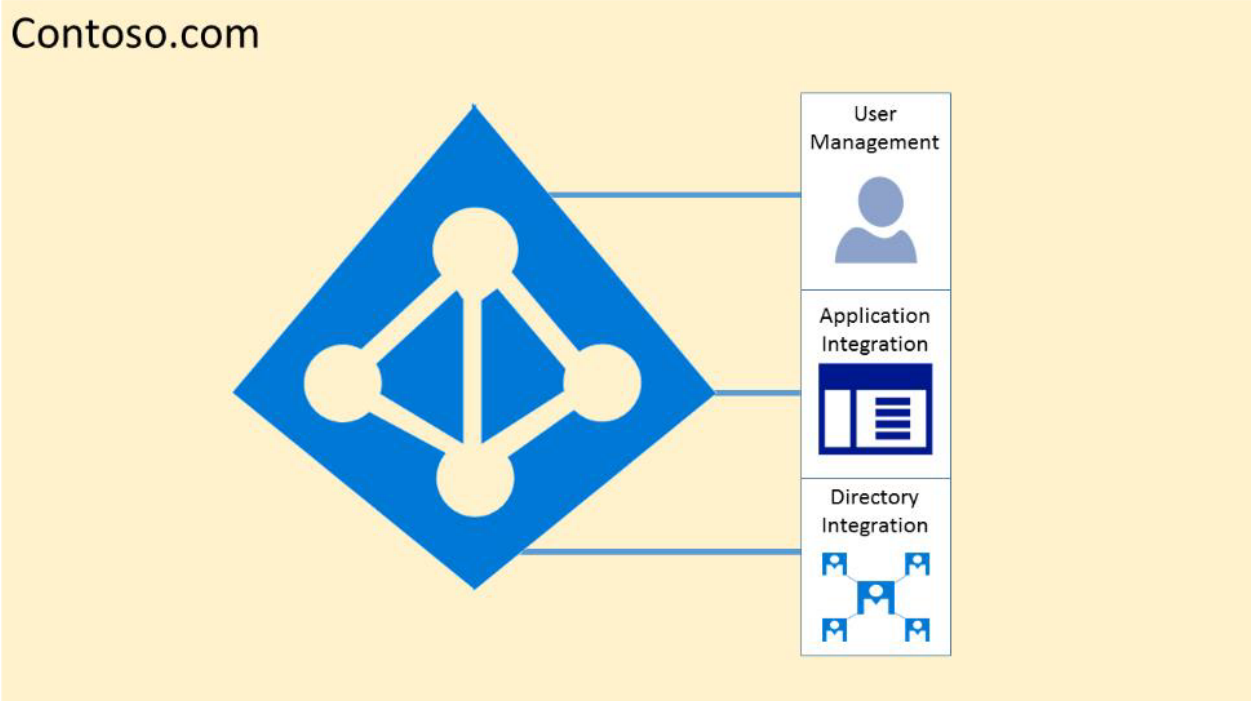
Azure tenancy (Azure Subscription) refers to a “customer/billing” relationship and a unique tenant in Microsoft Entra ID. Tenant level isolation in Microsoft Azure is achieved using Microsoft Entra ID and Azure role-based access control offered by it. Each Azure subscription is associated with one Microsoft Entra directory.

Users, groups, and applications from that directory can manage resources in the Azure subscription. You can assign these access rights using the Azure portal, Azure command-line tools, and Azure Management APIs. A Microsoft Entra tenant is logically isolated using security boundaries so that no customer can access or compromise co-tenants, either maliciously or accidentally. Microsoft Entra ID runs on “bare metal” servers isolated on a segregated network segment, where host-level packet filtering and Windows Firewall block unwanted connections and traffic.



- Access to data in Microsoft Entra ID requires user authentication via a security token service (STS). Information on the user’s existence, enabled state, and role is used by the authorization system to determine whether the requested access to the target tenant is authorized for this user in this session.
- Tenants are discrete containers and there's no relationship between these.
- No access across tenants unless tenant admin grants it through federation or provisioning user accounts from other tenants.
- Physical access to servers that comprise the Microsoft Entra service, and direct access to Microsoft Entra ID’s back-end systems, is restricted.
- Microsoft Entra users have no access to physical assets or locations, and therefore it isn't possible for them to bypass the logical Azure RBAC policy checks stated following.

For diagnostics and maintenance needs, an operational model that employs a just-in-time privilege elevation system is required and used. Microsoft Entra Privileged Identity Management (PIM) introduces the concept of an eligible admin. Eligible admins should be users that need privileged access now and then, but not every day. The role is inactive until the user needs access, then they complete an activation process and become an active admin for a predetermined amount of time.



Microsoft Entra ID hosts each tenant in its own protected container, with policies and permissions to and within the container solely owned and managed by the tenant.

The concept of tenant containers is deeply ingrained in the directory service at all layers, from portals all the way to persistent storage.

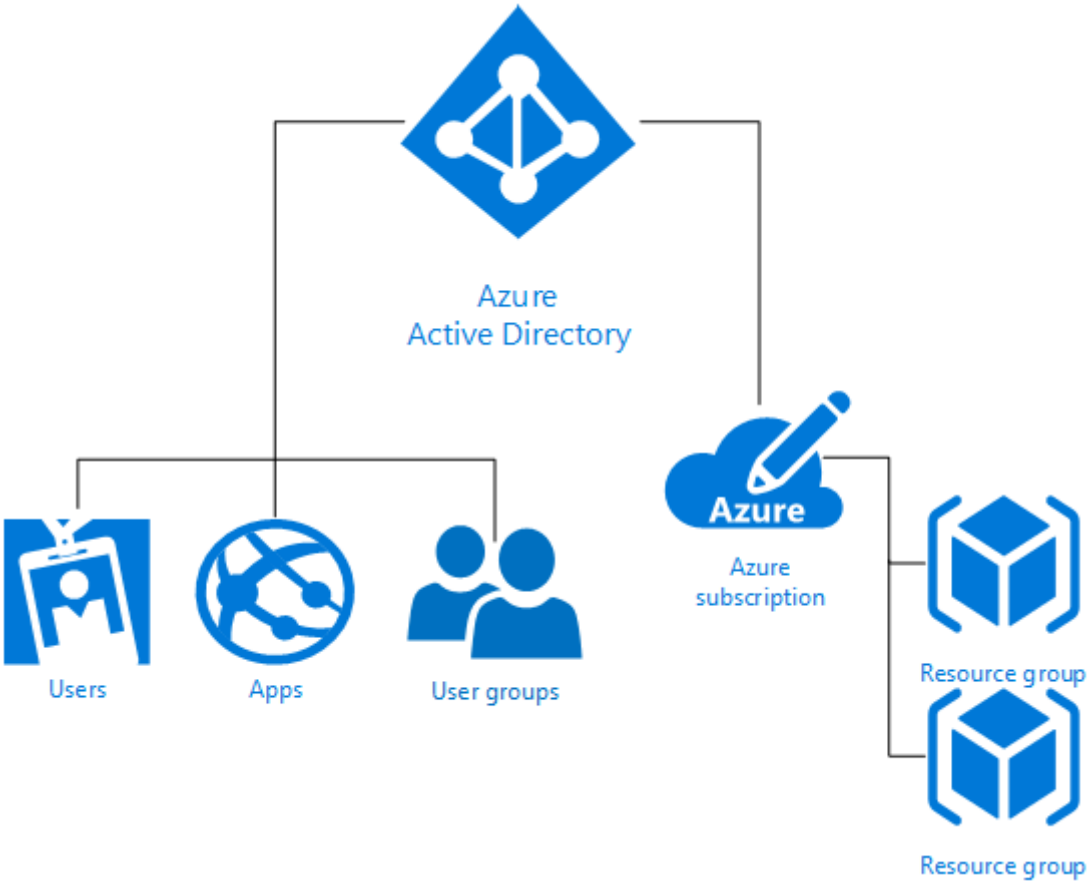
Even when metadata from multiple Microsoft Entra tenants is stored on the same physical disk, there's no relationship between the containers other than what is defined by the directory service, which in turn is dictated by the tenant administrator.

Azure role-based access control (Azure RBAC)

Azure role-based access control (Azure RBAC) helps you to share various components available within an Azure subscription by providing fine-grained access management for Azure. Azure RBAC enables you to segregate duties within your organization and grant access based on what users need to perform their jobs. Instead of giving everybody unrestricted permissions in Azure subscription or resources, you can allow only certain actions.

Azure RBAC has three basic roles that apply to all resource types:

- **Owner** has full access to all resources including the right to delegate access to others.
- **Contributor** can create and manage all types of Azure resources but can't grant access to others.
- **Reader** can view existing Azure resources.



The rest of the Azure roles in Azure allow management of specific Azure resources. For example, the Virtual Machine Contributor role allows the user to create and manage virtual machines. It doesn't give them access to the Azure Virtual Network or the subnet that the virtual machine connects to.

Azure built-in roles list the roles available in Azure. It specifies the operations and scope that each built-in role grants to users. If you're looking to define your own roles for even more control, see how to build Custom roles in Azure RBAC.

Some other capabilities for Microsoft Entra ID include:

- Microsoft Entra ID enables SSO to SaaS applications, regardless of where they're hosted. Some applications are federated with Microsoft Entra ID, and others use password SSO. Federated applications can also support user provisioning and password vaulting.

- Access to data in Azure Storage is controlled via authentication. Each storage account has a primary key (storage account key, or SAK) and a secondary secret key (the shared access signature, or SAS).

- Microsoft Entra ID provides Identity as a Service through federation by using Active Directory Federation Services, synchronization, and replication with on-premises directories.

- Microsoft Entra multifactor authentication requires users to verify sign-ins by using a mobile app, phone call, or text message. It can be used with Microsoft Entra ID to help secure on-premises resources with the Multi-Factor Authentication Server, and also with custom applications and directories using the SDK.

- Microsoft Entra Domain Services lets you join Azure virtual machines to an Active Directory domain without deploying domain controllers. You can sign in to these virtual machines with your corporate Active Directory credentials and administer domain-joined virtual machines by using Group Policy to enforce security baselines on all your Azure virtual machines.

- Azure Active Directory B2C provides a highly available global-identity management service for consumer-facing applications that scales to hundreds of millions of identities. It can be integrated across mobile and web platforms. Your consumers can sign in to all your applications through customizable experiences by using their existing social accounts or by creating credentials.

Isolation from Microsoft Administrators & Data Deletion

Microsoft takes strong measures to protect your data from inappropriate access or use by unauthorized persons. These operational processes and controls are backed by the Online Services Terms, which offer contractual commitments that govern access to your data.

- Microsoft engineers don't have default access to your data in the cloud. Instead, they're granted access, under management oversight, only when necessary. That access is carefully controlled and logged, and revoked when it's no longer needed.
- Microsoft may hire other companies to provide limited services on its behalf. Subcontractors may access customer data only to deliver the services for which, we have hired them to provide, and they're prohibited from using it for any other purpose. Further, they're contractually bound to maintain the confidentiality of our customers' information.

Business services with audited certifications such as ISO/IEC 27001 are regularly verified by Microsoft and accredited audit firms, which perform sample audits to attest that

access, only for legitimate business purposes. You can always access your own customer data at any time and for any reason.

If you delete any data, Microsoft Azure deletes the data, including any cached or backup copies. For in-scope services, that deletion will occur within 90 days after the end of the retention period. (In-scope services are defined in the Data Processing Terms section of our Online Services Terms .)

If a disk drive used for storage suffers a hardware failure, it's securely erased or destroyed before Microsoft returns it to the manufacturer for replacement or repair. The data on the drive is overwritten to ensure that the data can't be recovered by any means.

Compute Isolation

Microsoft Azure provides various cloud-based computing services that include a wide selection of compute instances & services that can scale up and down automatically to meet the needs of your application or enterprise. These compute instance and service offer isolation at multiple levels to secure data without sacrificing the flexibility in configuration that customers demand.

Isolated Virtual Machine Sizes

Azure Compute offers virtual machine sizes that are Isolated to a specific hardware type and dedicated to a single customer. The Isolated sizes live and operate on specific hardware generation and will be deprecated when the hardware generation is retired or new hardware generation is available.

Isolated virtual machine sizes are best suited for workloads that require a high degree of isolation from other customers' workloads. This is sometimes required to meet compliance and regulatory requirements. Utilizing an isolated size guarantees that your virtual machine is the only one running on that specific server instance.

Additionally, as the Isolated size VMs are large, customers may choose to subdivide the resources of these VMs by using Azure support for nested virtual machines .

The current Isolated virtual machine offerings include:

- Standard_E80ids_v4
- Standard_E80is_v4
- Standard_E104i_v5
- Standard_E104is_v5

- Standard_E104id_v5
- Standard_E104ids_v5
- Standard_M192is_v2
- Standard_M192ims_v2
- Standard_M192ids_v2
- Standard_M192idms_v2
- Standard_F72s_v2
- Standard_M128ms

> ⓘ **Note**
>
> Isolated VM Sizes have a limited lifespan due to hardware deprecation.

Deprecation of Isolated VM Sizes

Isolated VM sizes have a hardware limited lifespan. Azure issues reminders 12 months in advance of the official deprecation date of the sizes and provides an updated isolated offering for your consideration. The following sizes have retirement announced.

Size	Isolation Retirement Date
Standard_DS15_v2	May 15, 2021
Standard_D15_v2	May 15, 2021
Standard_G5	February 15, 2022
Standard_GS5	February 15, 2022
Standard_E64i_v3	February 15, 2022
Standard_E64is_v3	February 15, 2022

FAQ

Q: Is the size going to get retired or only its "isolation" feature?

A: Any size that is published as isolated but have no "i" in the name, the isolation feature of the VM sizes is being retired unless communicated differently. Sizes with "i" in the name will be deprecated.

Q: Is there a downtime when my vm lands on a nonisolated hardware?

A: For VM sizes, where only isolation is deprecating but not the size, no action is needed and there will be no downtime. On contrary if isolation is required, announcement includes the recommended replacement size. Selecting the replacement size requires customers to resize their VMs.

Q: Is there any cost delta for moving to a nonisolated virtual machine?

A: No

Q: When are the other isolated sizes going to retire?

A: We provide reminders 12 months in advance of the official deprecation of the isolated size. Our latest announcement includes isolation feature retirement of Standard_G5, Standard_GS5, Standard_E64i_v3 and Standard_E64i_v3.

Q: I'm an Azure Service Fabric Customer relying on the Silver or Gold Durability Tiers. Does this change impact me?

A: No. The guarantees provided by Service Fabric's Durability Tiers will continue to function even after this change. If you require physical hardware isolation for other reasons, you may still need to take one of the actions described above.

Q: What are the milestones for D15_v2 or DS15_v2 isolation retirement?

A:

Date	Action
May 15, 2020[1]	D/DS15_v2 isolation retirement announcement
May 15, 2021	D/DS15_v2 isolation guarantee removed

[1] Existing customer using these sizes will receive an announcement email with detailed instructions on the next steps.

Q: What are the milestones for G5, Gs5, E64i_v3 and E64is_v3 isolation retirement?

A:

Date	Action
Feb 15, 2021[1]	G5/GS5/E64i_v3/E64is_v3 isolation retirement announcement
Feb 28, 2022	G5/GS5/E64i_v3/E64is_v3 isolation guarantee removed

[1] Existing customer using these sizes will receive an announcement email with detailed instructions on the next steps.

Next steps

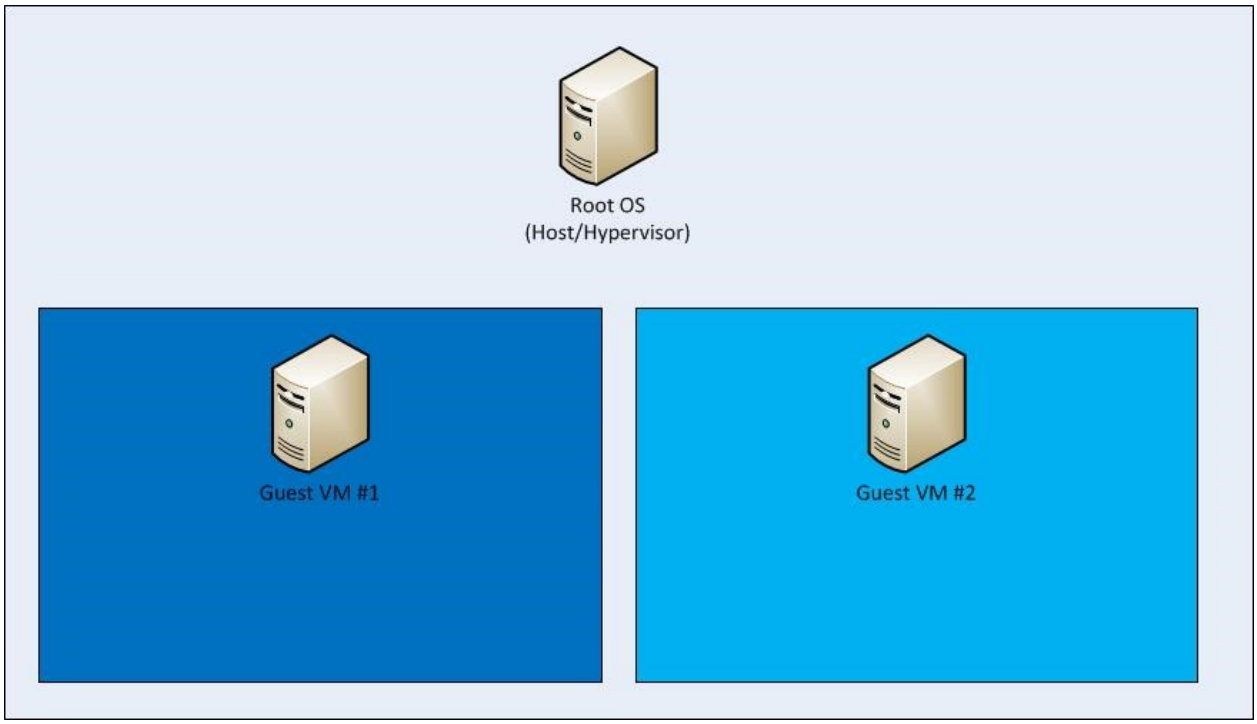
Customers can also choose to further subdivide the resources of these Isolated virtual machines by using Azure support for nested virtual machines.

Dedicated hosts

In addition to the isolated hosts described in the preceding section, Azure also offers dedicated hosts. Dedicated hosts in Azure is a service that provides physical servers that can host one or more virtual machines, and which are dedicated to a single Azure subscription. Dedicated hosts provide hardware isolation at the physical server level. No other VMs will be placed on your hosts. Dedicated hosts are deployed in the same datacenters and share the same network and underlying storage infrastructure as other, non-isolated hosts. For more information, see the detailed overview of Azure dedicated hosts.

Hyper-V & Root OS Isolation Between Root VM & Guest VMs

Azure's compute platform is based on machine virtualization—meaning that all customer code executes in a Hyper-V virtual machine. On each Azure node (or network endpoint), there's a Hypervisor that runs directly over the hardware and divides a node into a variable number of Guest Virtual Machines (VMs).



Each node also has one special Root VM, which runs the Host OS. A critical boundary is the isolation of the root VM from the guest VMs and the guest VMs from one another, managed by the hypervisor and the root OS. The hypervisor/root OS pairing leverages Microsoft's decades of operating system security experience, and more recent learning from Microsoft's Hyper-V, to provide strong isolation of guest VMs.

The Azure platform uses a virtualized environment. User instances operate as standalone virtual machines that don't have access to a physical host server.

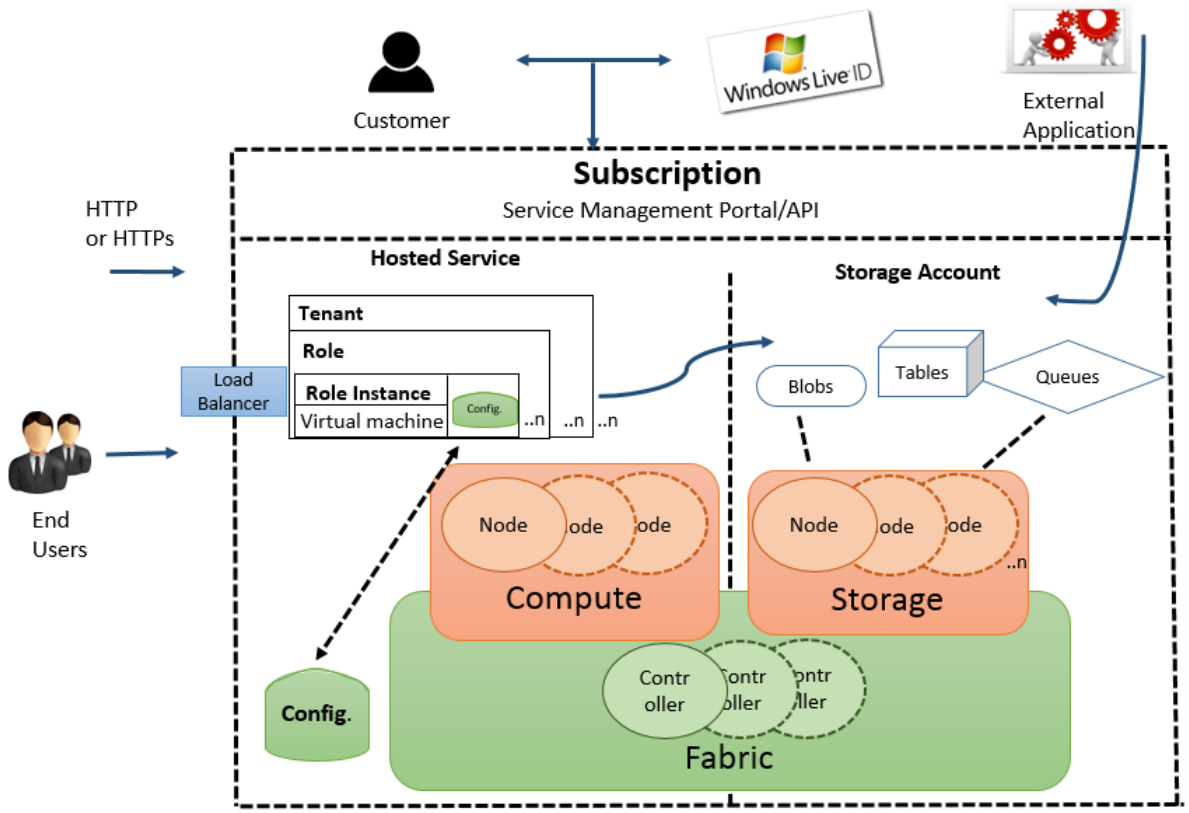
The Azure hypervisor acts like a micro-kernel and passes all hardware access requests from guest virtual machines to the host for processing by using a shared-memory interface called VM Bus. This prevents users from obtaining raw read/write/execute access to the system and mitigates the risk of sharing system resources.

Advanced VM placement algorithm & protection from side channel attacks

Any cross-VM attack involves two steps: placing an adversary-controlled VM on the same host as one of the victim VMs, and then breaching the isolation boundary to either steal sensitive victim information or affect its performance for greed or vandalism. Microsoft Azure provides protection at both steps by using an advanced VM placement algorithm and protection from all known side channel attacks including noisy neighbor VMs.

The Azure Fabric Controller

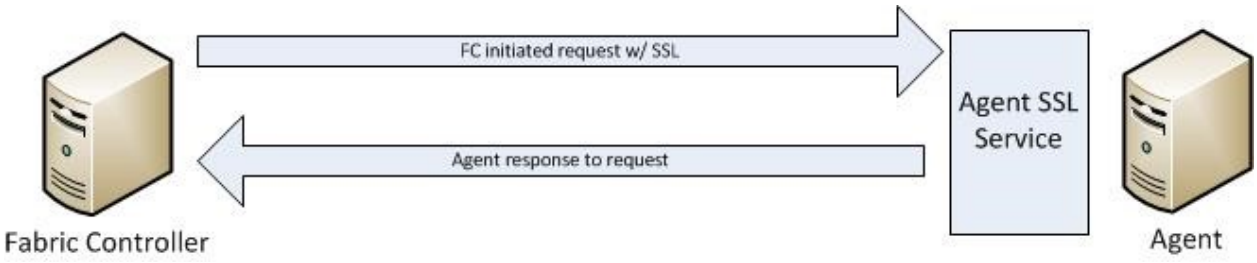
The Azure Fabric Controller is responsible for allocating infrastructure resources to tenant workloads, and it manages unidirectional communications from the host to virtual machines. The VM placing algorithm of the Azure fabric controller is highly sophisticated and nearly impossible to predict as physical host level.



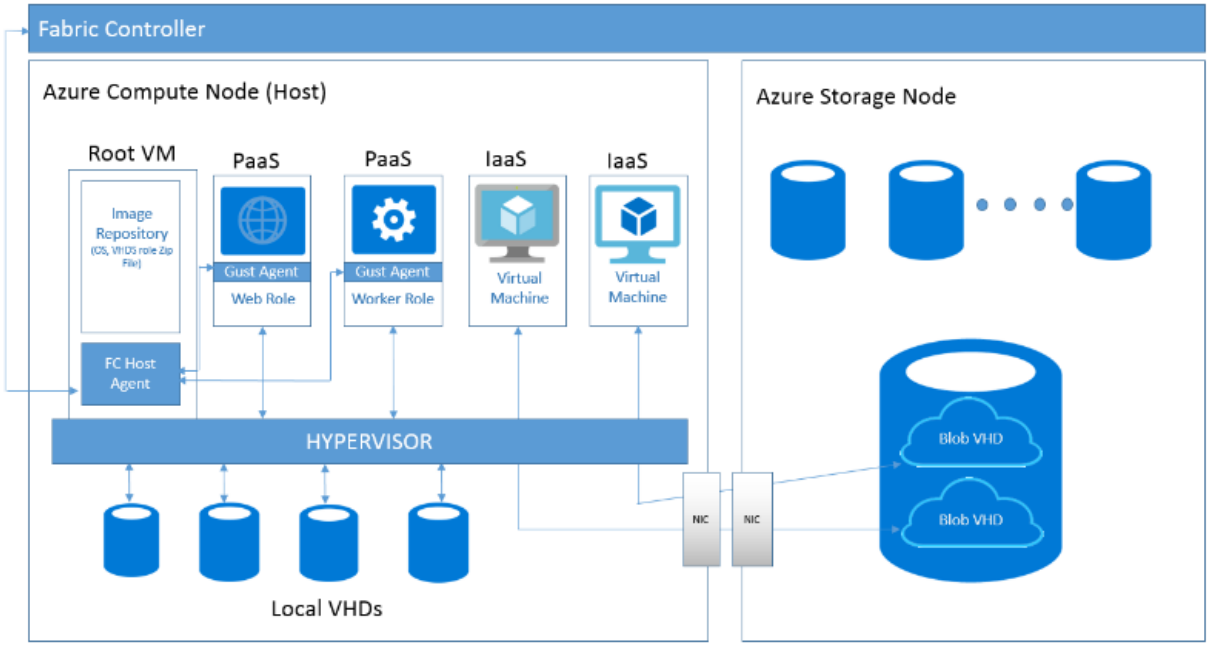
The Azure hypervisor enforces memory and process separation between virtual machines, and it securely routes network traffic to guest OS tenants. This eliminates possibility of and side channel attack at VM level.

In Azure, the root VM is special: it runs a hardened operating system called the root OS that hosts a fabric agent (FA). FAs are used in turn to manage guest agents (GA) within guest operating systems on customer VMs. FAs also manage storage nodes.

The collection of Azure hypervisor, root OS/FA, and customer VMs/GAs comprises a compute node. FAs are managed by a fabric controller (FC), which exists outside of compute and storage nodes (compute and storage clusters are managed by separate FCs). If a customer updates their application's configuration file while it's running, the FC communicates with the FA, which then contacts GAs, which notify the application of the configuration change. In the event of a hardware failure, the FC will automatically find available hardware and restart the VM there.



Communication from a Fabric Controller to an agent is unidirectional. The agent implements an SSL-protected service that only responds to requests from the controller. It cannot initiate connections to the controller or other privileged internal nodes. The FC treats all responses as if they were untrusted.



Isolation extends from the Root VM from Guest VMs, and the Guest VMs from one another. Compute nodes are also isolated from storage nodes for increased protection.

The hypervisor and the host OS provide network packet - filters to help assure that untrusted virtual machines cannot generate spoofed traffic or receive traffic not addressed to them, direct traffic to protected infrastructure endpoints, or send/receive inappropriate broadcast traffic.

Additional Rules Configured by Fabric Controller Agent to Isolate VM

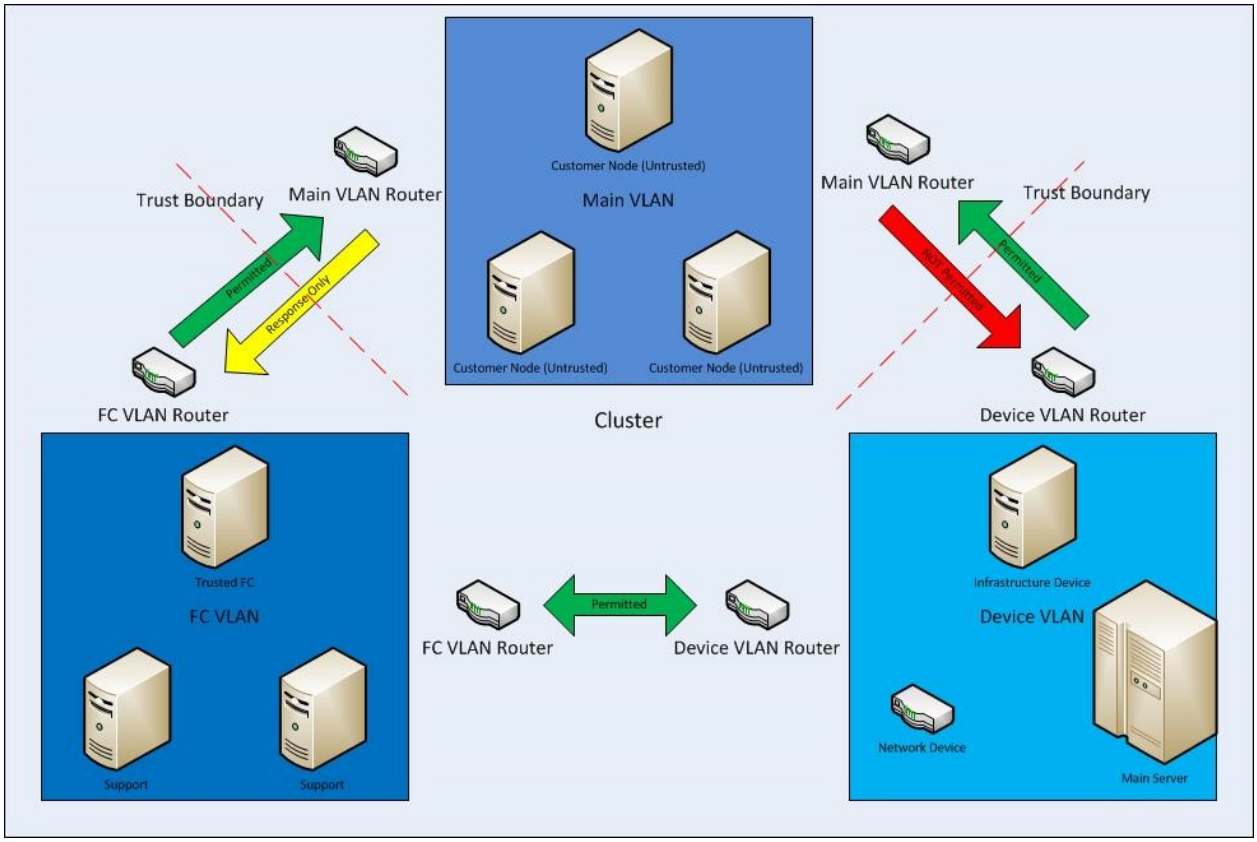
By default, all traffic is blocked when a virtual machine is created, and then the fabric controller agent configures the packet filter to add rules and exceptions to allow authorized traffic.

There are two categories of rules that are programmed:

- **Machine configuration or infrastructure rules:** By default, all communication is blocked. There are exceptions to allow a virtual machine to send and receive DHCP and DNS traffic. Virtual machines can also send traffic to the “public” internet and send traffic to other virtual machines within the same Azure Virtual Network and the OS activation server. The virtual machines’ list of allowed outgoing destinations doesn't include Azure router subnets, Azure management, and other Microsoft properties.
- **Role configuration file:** This defines the inbound Access Control Lists (ACLs) based on the tenant's service model.

VLAN Isolation

There are three VLANs in each cluster:



- The main VLAN – interconnects untrusted customer nodes
- The FC VLAN – contains trusted FCs and supporting systems
- The device VLAN – contains trusted network and other infrastructure devices

Communication is permitted from the FC VLAN to the main VLAN, but cannot be initiated from the main VLAN to the FC VLAN. Communication is also blocked from the main VLAN to the device VLAN. This assures that even if a node running customer code is compromised, it cannot attack nodes on either the FC or device VLANs.

Storage Isolation

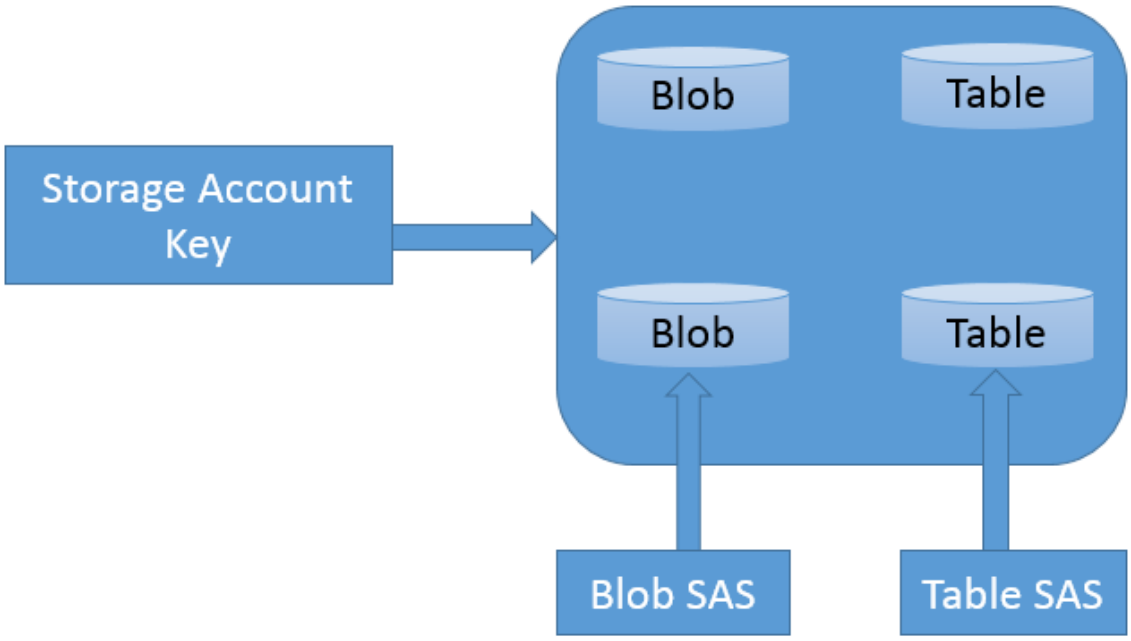
Logical Isolation Between Compute and Storage

As part of its fundamental design, Microsoft Azure separates VM-based computation from storage. This separation enables computation and storage to scale independently, making it easier to provide multi-tenancy and isolation.

Therefore, Azure Storage runs on separate hardware with no network connectivity to Azure Compute except logically. This means that when a virtual disk is created, disk space isn't allocated for its entire capacity. Instead, a table is created that maps addresses on the virtual disk to areas on the physical disk and that table is initially empty. **The first time a customer writes data on the virtual disk, space on the physical disk is allocated, and a pointer to it's placed in the table.**

Isolation Using Storage Access control

Access Control in Azure Storage has a simple access control model. Each Azure subscription can create one or more Storage Accounts. Each Storage Account has a single secret key that's used to control access to all data in that Storage Account.



Access to Azure Storage data (including Tables) can be controlled through a SAS (Shared Access Signature) token, which grants scoped access. The SAS is created through a query template (URL), signed with the SAK (Storage Account Key). That signed URL can be given to another process (that is, delegated), which can then fill in the details of the query and make the request of the storage service. A SAS enables you to grant time-based access to clients without revealing the storage account's secret key.

The SAS means that we can grant a client limited permissions, to objects in our storage account for a specified period of time and with a specified set of permissions. We can grant these limited permissions without having to share your account access keys.

IP Level Storage Isolation

You can establish firewalls and define an IP address range for your trusted clients. With an IP address range, only clients that have an IP address within the defined range can connect to Azure Storage.

IP storage data can be protected from unauthorized users via a networking mechanism that's used to allocate a dedicated or dedicated tunnel of traffic to IP storage.

Encryption

Azure offers the following types of Encryption to protect data:

- Encryption in transit
- Encryption at rest

Encryption in Transit

Encryption in transit is a mechanism of protecting data when it's transmitted across networks. With Azure Storage, you can secure data using:

- Transport-level encryption, such as HTTPS when you transfer data into or out of Azure Storage.
- Wire encryption, such as SMB 3.0 encryption for Azure File shares.
- Client-side encryption, to encrypt the data before it's transferred into storage and to decrypt the data after it's transferred out of storage.

Encryption at Rest

For many organizations, data encryption at rest is a mandatory step towards data privacy, compliance, and data sovereignty. There are three Azure features that provide

encryption of data that's "at rest":

- Storage Service Encryption allows you to request that the storage service automatically encrypt data when writing it to Azure Storage.
- Client-side Encryption also provides the feature of encryption at rest.
- Azure Disk Encryption for Linux VMs and Azure Disk Encryption for Windows VMs.

For more information, see Overview of managed disk encryption options.

Azure Disk Encryption

Azure Disk Encryption for Linux VMs and Azure Disk Encryption for Windows VMs help you address organizational security and compliance requirements by encrypting your VM disks (including boot and data disks) with keys and policies you control in Azure Key Vault.

The Disk Encryption solution for Windows is based on Microsoft BitLocker Drive Encryption, and the Linux solution is based on dm-crypt.

The solution supports the following scenarios for IaaS VMs when they're enabled in Microsoft Azure:

- Integration with Azure Key Vault
- Standard tier VMs: A, D, DS, G, GS, and so forth, series IaaS VMs
- Enabling encryption on Windows and Linux IaaS VMs
- Disabling encryption on OS and data drives for Windows IaaS VMs
- Disabling encryption on data drives for Linux IaaS VMs
- Enabling encryption on IaaS VMs that are running Windows client OS
- Enabling encryption on volumes with mount paths
- Enabling encryption on Linux VMs that are configured with disk striping (RAID) by using mdadm
- Enabling encryption on Linux VMs by using LVM(Logical Volume Manager) for data disks
- Enabling encryption on Windows VMs that are configured by using storage spaces
- All Azure public regions are supported

The solution doesn't support the following scenarios, features, and technology in the release:

- Basic tier IaaS VMs
- Disabling encryption on an OS drive for Linux IaaS VMs
- IaaS VMs that are created by using the classic VM creation method
- Integration with your on-premises Key Management Service

- Azure Files (shared file system), Network File System (NFS), dynamic volumes, and Windows VMs that are configured with software-based RAID systems

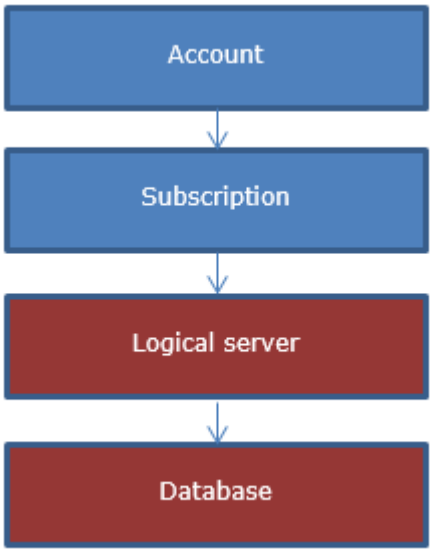
SQL Database Isolation

SQL Database is a relational database service in the Microsoft cloud based on the market-leading Microsoft SQL Server engine and capable of handling mission-critical workloads. SQL Database offers predictable data isolation at account level, geography / region based and based on networking— all with near-zero administration.

SQL Database Application Model

Microsoft SQL Database is a cloud-based relational database service built on SQL Server technologies. It provides a highly available, scalable, multi-tenant database service hosted by Microsoft in cloud.

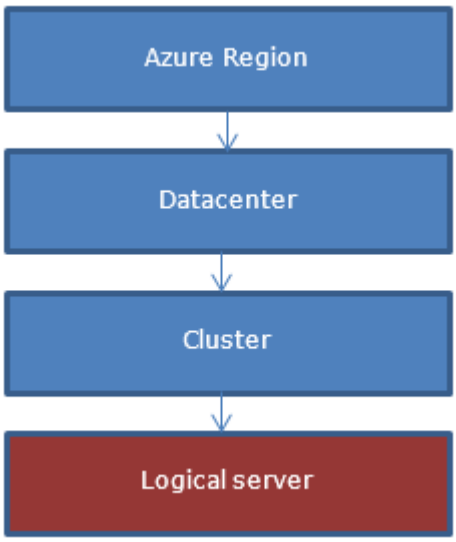
From an application perspective, SQL Database provides the following hierarchy: Each level has one-to-many containment of levels below.



The account and subscription are Microsoft Azure platform concepts to associate billing and management.

Logical SQL servers and databases are SQL Database-specific concepts and are managed by using SQL Database, provided OData and TSQL interfaces or via the Azure portal.

Servers in SQL Database aren't physical or VM instances, instead they 're collections of databases, sharing management and security policies, which are stored in so called “logical master” database.



Logical master databases include:

- SQL logins used to connect to the server
- Firewall rules

Billing and usage-related information for databases from the same server aren't guaranteed to be on the same physical instance in the cluster, instead applications must provide the target database name when connecting.

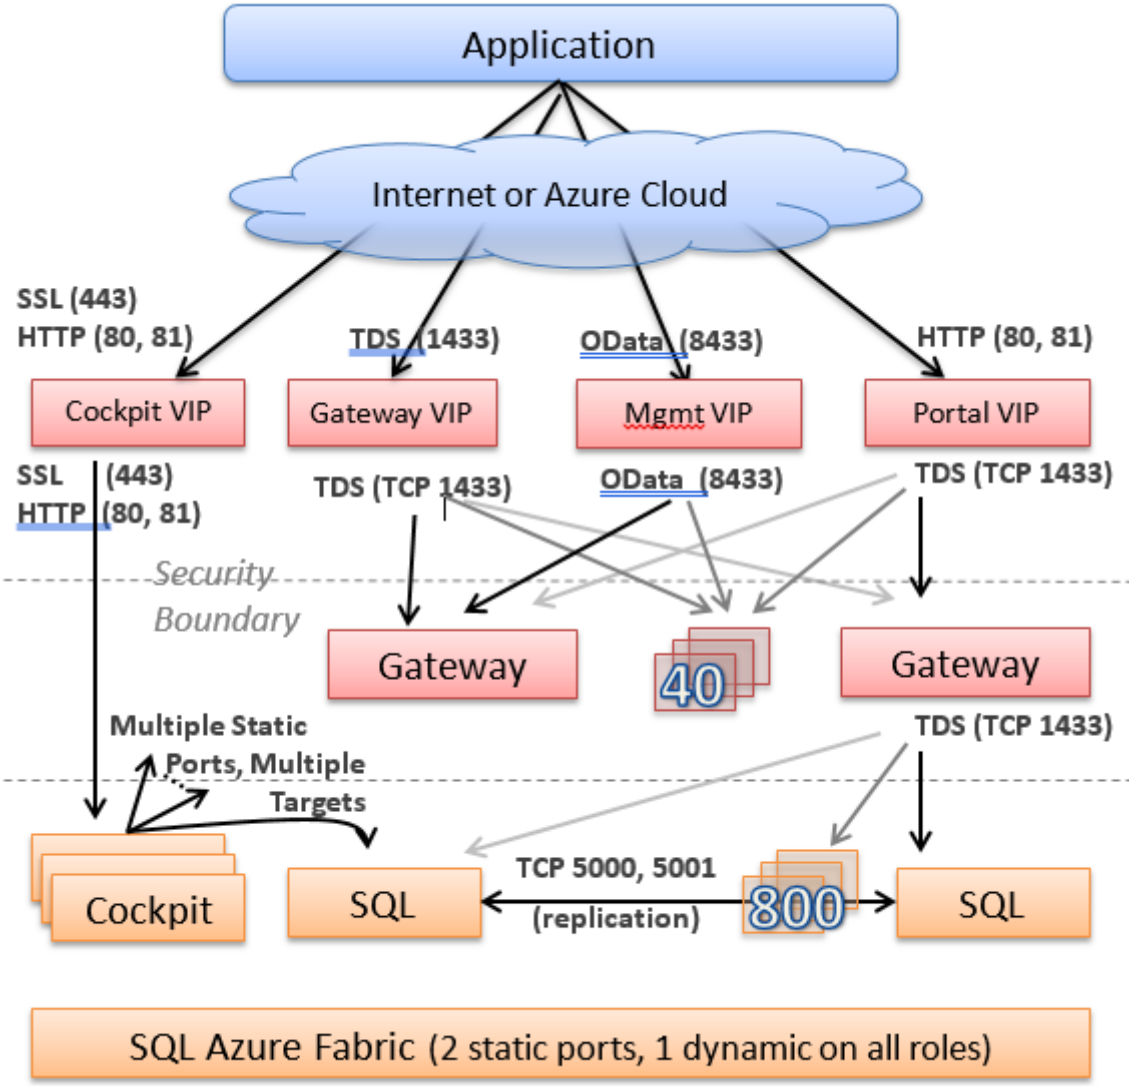
From a customer perspective, a server is created in a geo-graphical region while the actual creation of the server happens in one of the clusters in the region.

Isolation through Network Topology

When a server is created and its DNS name is registered, the DNS name points to the so called “Gateway VIP” address in the specific data center where the server was placed.

Behind the VIP (virtual IP address), we have a collection of stateless gateway services. In general, gateways get involved when there's coordination needed between multiple data sources (master database, user database, etc.). Gateway services implement the following:

- **TDS connection proxying.** This includes locating user database in the backend cluster, implementing the login sequence and then forwarding the TDS packets to the backend and back.
- **Database management.** This includes implementing a collection of workflows to do CREATE/ALTER/DROP database operations. The database operations can be invoked by either sniffing TDS packets or explicit OData APIs.
- CREATE/ALTER/DROP login/user operations
- Server management operations via OData API



The tier behind the gateways is called “back-end”. This is where all the data is stored in a highly available fashion. Each piece of data is said to belong to a “partition” or “failover unit”, each of them having at least three replicas. Replicas are stored and replicated by SQL Server engine and managed by a failover system often referred to as “fabric”.

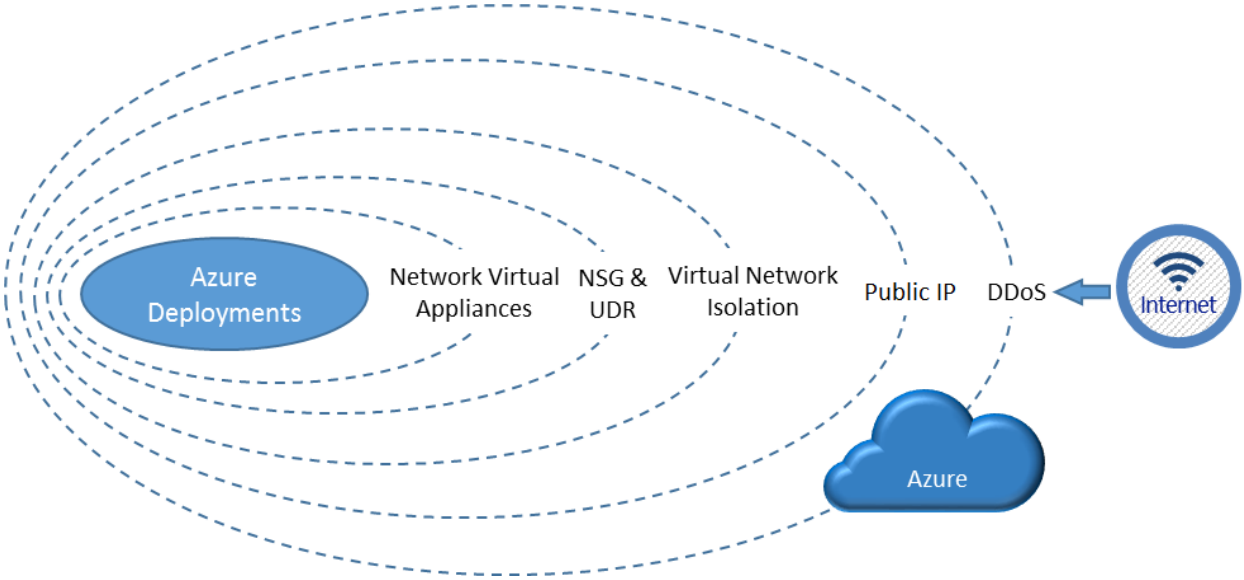
Generally, the back-end system doesn't communicate outbound to other systems as a security precaution. This is reserved to the systems in the front-end (gateway) tier. The gateway tier machines have limited privileges on the back-end machines to minimize the attack surface as a defense-in-depth mechanism.

Isolation by Machine Function and Access

SQL Database is composed of services running on different machine functions. SQL Database is divided into “backend” Cloud Database and “front-end” (Gateway/Management) environments, with the general principle of traffic only going into back-end and not out. The front-end environment can communicate to the outside world of other services and in general, has only limited permissions in the back-end (enough to call the entry points it needs to invoke).

Networking Isolation

Azure deployment has multiple layers of network isolation. The following diagram shows various layers of network isolation Azure provides to customers. These layers are both native in the Azure platform itself and customer-defined features. Inbound from the Internet, Azure DDoS provides isolation against large-scale attacks against Azure. The next layer of isolation is customer-defined public IP addresses (endpoints), which are used to determine which traffic can pass through the cloud service to the virtual network. Native Azure virtual network isolation ensures complete isolation from all other networks, and that traffic only flows through user configured paths and methods. These paths and methods are the next layer, where NSGs, UDR, and network virtual appliances can be used to create isolation boundaries to protect the application deployments in the protected network.



Traffic isolation: A virtual network is the traffic isolation boundary on the Azure platform. Virtual machines (VMs) in one virtual network cannot communicate directly to VMs in a different virtual network, even if both virtual networks are created by the same customer. Isolation is a critical property that ensures customer VMs and communication remains private within a virtual network.

Subnet offers an additional layer of isolation with in virtual network based on IP range. IP addresses in the virtual network, you can divide a virtual network into multiple subnets for organization and security. VMs and PaaS role instances deployed to subnets (same or different) within a VNet can communicate with each other without any extra configuration. You can also configure network security group (NSGs) to allow or deny network traffic to a VM instance based on rules configured in access control list (ACL) of NSG. NSGs can be associated with either subnets or individual VM instances within that subnet. When an NSG is associated with a subnet, the ACL rules apply to all the VM instances in that subnet.

Next Steps

- Learn about Network Isolation Options for Machines in Windows Azure Virtual Networks. This includes the classic front-end and back-end scenario where machines in a particular back-end network or subnetwork may only allow certain clients or other computers to connect to a particular endpoint based on an allowlist of IP addresses.
- Learn about virtual machine isolation in Azure. Azure Compute offers virtual machine sizes that are isolated to a specific hardware type and dedicated to a single customer.

Microsoft 365 isolation controls

Article • 05/25/2023

Microsoft continuously works to ensure that the multi-tenant architecture of Microsoft 365 supports enterprise-level security, confidentiality, privacy, integrity, local, international, and availability standards. The scale and the scope of services provided by Microsoft make it difficult and non-economical to manage Microsoft 365 with significant human interaction. Microsoft 365 services are provided through globally distributed data centers, each highly automated with few operations requiring a human touch or any access to customer content. Our staff supports these services and data centers using automated tools and highly secure remote access.

Microsoft 365 is composed of multiple services that provide important business functionality and contribute to the entire Microsoft 365 experience. Each of these services is self-contained and designed to integrate with one another. Microsoft 365 is designed with the following principles:

- Service-oriented architecture: designing and developing software in the form of interoperable services providing well-defined business functionality.
- Operational security assurance: a framework that incorporates the knowledge gained through various capabilities that are unique to Microsoft, including the Microsoft Security Development Lifecycle, the Microsoft Security Response Center, and deep awareness of the cybersecurity threat landscape.

Microsoft 365 services inter-operate with each other but are designed and implemented so they can be deployed and operated as autonomous services, independent of each other. Microsoft segregates duties and areas of responsibility for Microsoft 365 to reduce opportunities for unauthorized or unintentional modification or misuse of the organization's assets. Microsoft 365 teams have defined roles as part of a comprehensive role-based access control mechanism.

Tenant isolation

One of the primary benefits of cloud computing is concept of a shared, common infrastructure across numerous customers simultaneously, leading to economies of scale. Microsoft works continuously to ensure that the multi-tenant architectures of our cloud services support enterprise-level security, confidentiality, privacy, integrity, and availability standards.

Microsoft cloud services were designed with the assumption that all tenants are potentially hostile to all other tenants, and we have implemented security measures to

prevent the actions of one tenant from affecting the security or service of another tenant or accessing the content of another tenant.

The two primary goals of maintaining tenant isolation in a multi-tenant environment are:

- Preventing leakage of, or unauthorized access to, customer content across tenants; and
- Preventing the actions of one tenant from adversely affecting the service for another tenant

Multiple forms of protection have been implemented throughout Microsoft 365 to prevent customers from compromising Microsoft 365 services or applications or gaining unauthorized access to the information of other tenants or the Microsoft 365 system itself, including:

- Logical isolation of customer content within each tenant for Microsoft 365 services is achieved through Microsoft Entra authorization and role-based access control.
- SharePoint Online provides data isolation mechanisms at the storage level.
- Microsoft uses rigorous physical security, background screening, and a multi-layered encryption strategy to protect the confidentiality and integrity of customer content. All Microsoft 365 datacenters have biometric access controls, with most requiring palm prints to gain physical access. In addition, all U.S.-based Microsoft employees are required to successfully complete a standard background check as part of the hiring process. For more information on the controls used for administrative access in Microsoft 365, see Microsoft 365 Account Management.
- Microsoft 365 uses service-side technologies that encrypt customer content at rest and in transit, including BitLocker, per-file encryption, Transport Layer Security (TLS) and Internet Protocol Security (IPsec). For specific details about encryption in Microsoft 365, see Data Encryption Technologies in Microsoft 365.

Together, the above-listed protections provide robust logical isolation controls that provide threat protection and mitigation equivalent to that provided by physical isolation alone.

Resources

- Isolation and access control in Microsoft Entra ID
- Monitoring and testing tenant boundaries
- Isolation and access control in Microsoft 365

Microsoft 365 Isolation and Access Control in Microsoft Entra ID

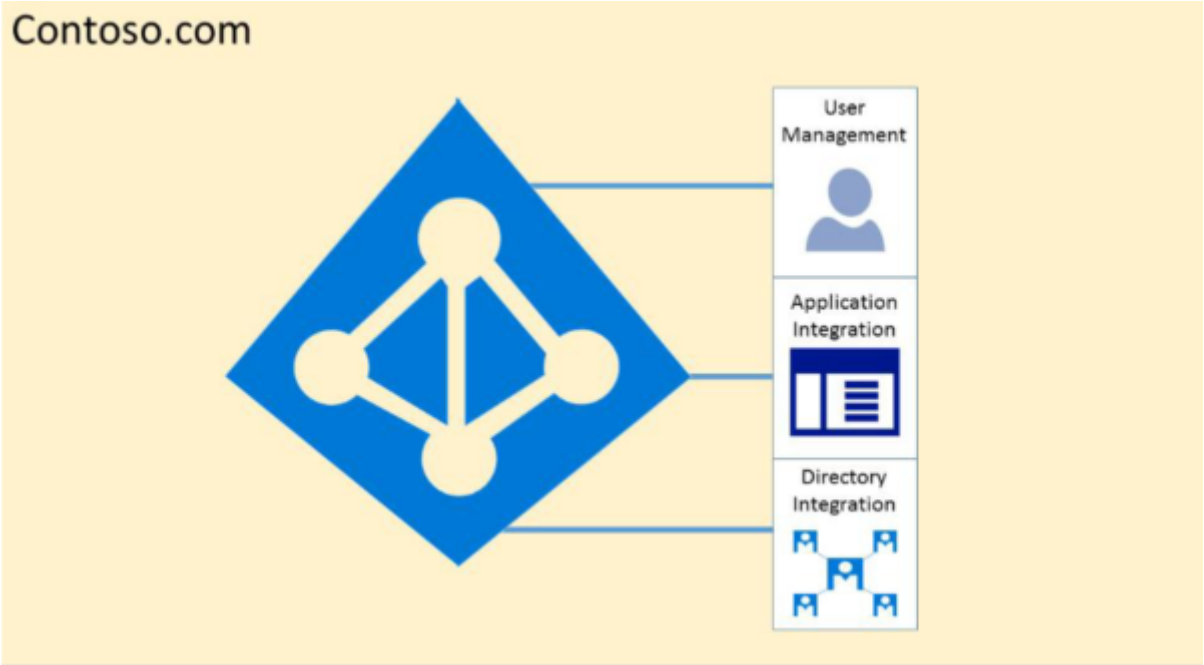
Article • 10/20/2023

Microsoft Entra ID was designed to host multiple tenants in a highly secure way through logical data isolation. Access to Microsoft Entra ID is gated by an authorization layer. Microsoft Entra ID isolates customers using tenant containers as security boundaries to safeguard a customer's content so that the content cannot be accessed or compromised by co-tenants. Three checks are performed by Microsoft Entra authorization layer:

- Is the principal enabled for access to Microsoft Entra tenant?
- Is the principal enabled for access to data in this tenant?
- Is the principal's role in this tenant authorized for the type of data access requested?

No application, user, server, or service can access Microsoft Entra ID without the proper authentication and token or certificate. Requests are rejected if they are not accompanied by proper credentials.

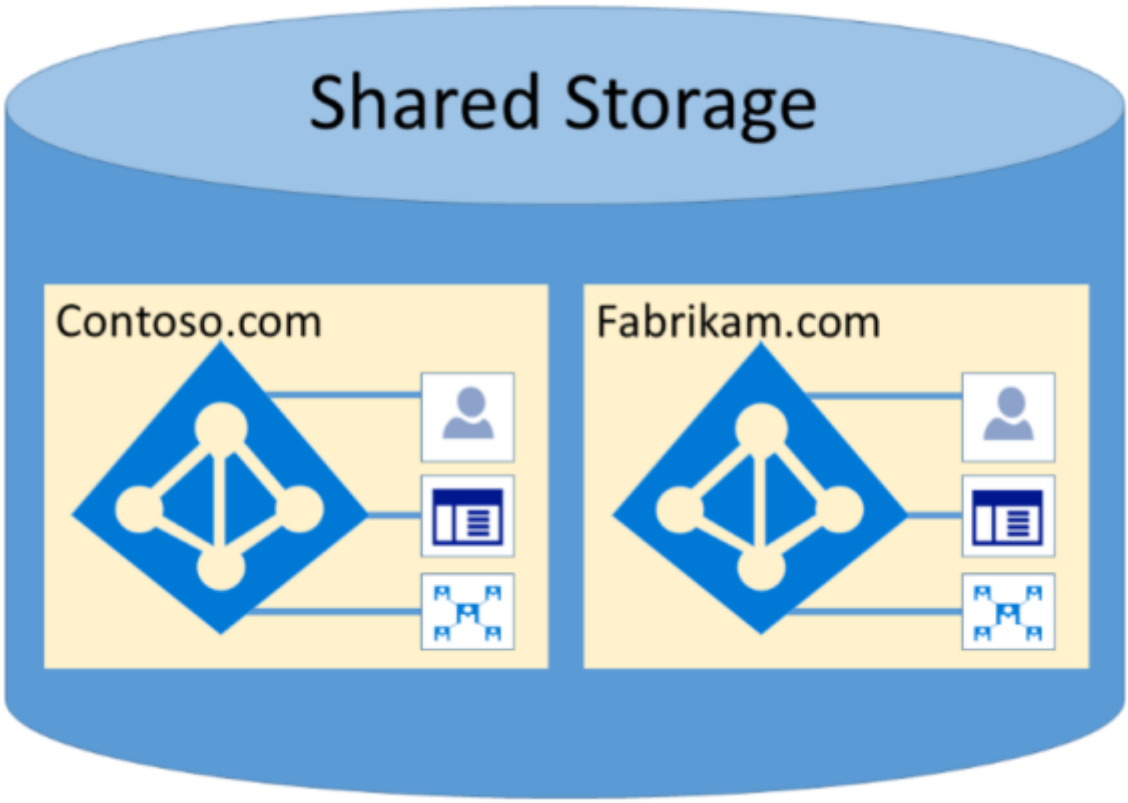
Effectively, Microsoft Entra ID hosts each tenant in its own protected container, with policies and permissions to and within the container solely owned and managed by the tenant.



The concept of tenant containers is deeply ingrained in the directory service at all layers, from portals all the way to persistent storage. Even when multiple Microsoft Entra tenant metadata is stored on the same physical disk, there is no relationship between the

containers other than what is defined by the directory service, which in turn is dictated by the tenant administrator. There can be no direct connections to Microsoft Entra storage from any requesting application or service without first going through the authorization layer.

In the example below, Contoso and Fabrikam both have separate, dedicated containers, and even though those containers may share some of the same underlying infrastructure, such as servers and storage, they remain separate and isolated from each other, and gated by layers of authorization and access control.



In addition, there are no application components that can execute from within Microsoft Entra ID, and it is not possible for one tenant to forcibly breach the integrity of another tenant, access encryption keys of another tenant, or read raw data from the server.

By default, Microsoft Entra disallows all operations issued by identities in other tenants. Each tenant is logically isolated within Microsoft Entra ID through claims-based access controls. Reads and writes of directory data are scoped to tenant containers, and gated by an internal abstraction layer and a role-based access control (RBAC) layer, which together enforce the tenant as the security boundary. Every directory data access request is processed by these layers and every access request in Microsoft 365 is policed by the logic above.

Microsoft Entra ID has North America, U.S. Government, European Union, Germany, and World Wide partitions. A tenant exists in a single partition, and partitions can contain

multiple tenants. Partition information is abstracted away from users. A given partition (including all the tenants within it) is replicated to multiple datacenters. The partition for a tenant is chosen based on properties of the tenant (e.g., the country code). Secrets and other sensitive information in each partition is encrypted with a dedicated key. The keys are generated automatically when a new partition is created.

Microsoft Entra system functionalities are a unique instance to each user session. In addition, Microsoft Entra ID uses encryption technologies to provide isolation of shared system resources at the network level to prevent unauthorized and unintended transfer of information.

Audit logging and monitoring overview

Article • 04/04/2023

How do Microsoft online services employ audit logging?

Microsoft online services employ audit logging to detect unauthorized activities and provide accountability for Microsoft personnel. Audit logs capture details about system configuration changes and access events, with details to identify who was responsible for the activity, when and where the activity took place, and what the outcome of the activity was. Automated log analysis supports near real-time detection of suspicious behavior. Potential incidents are escalated to the appropriate Microsoft security response team for further investigation.

Microsoft online services internal audit logging captures log data from various sources, such as:

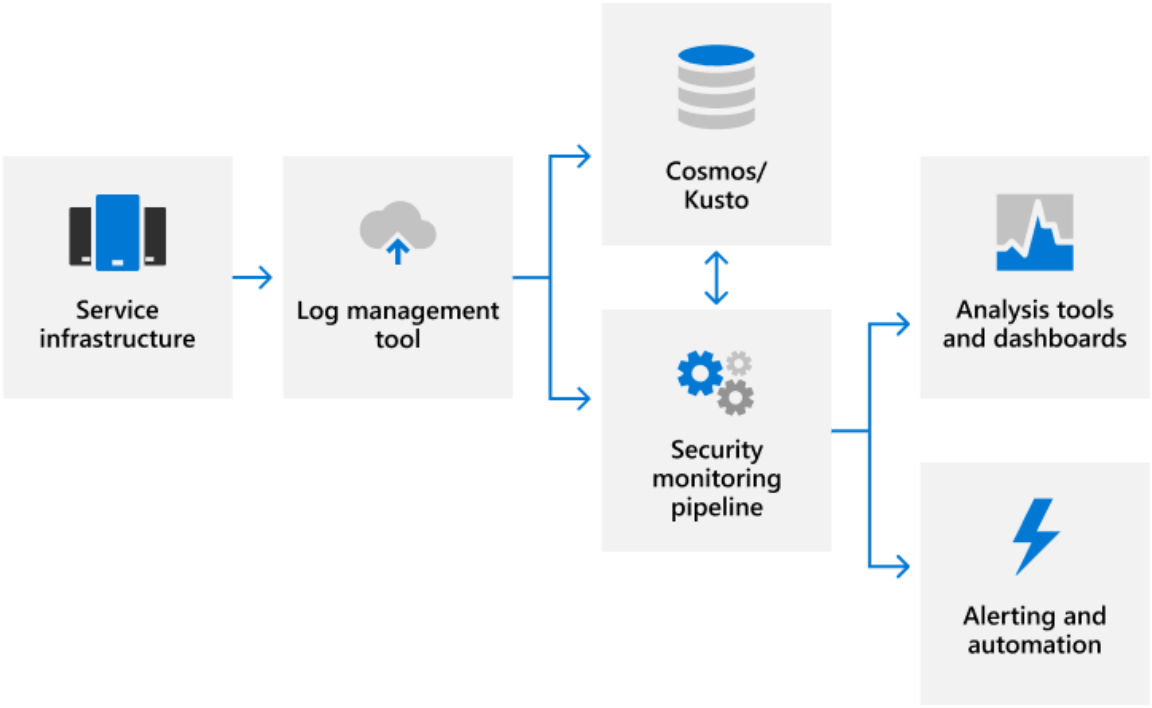
- Event logs
- AppLocker logs
- Performance data
- System Center data
- Call detail records
- Quality of experience data
- IIS Web Server logs
- SQL Server logs
- Syslog data
- Security audit logs

How do Microsoft online services centralize and report on audit logs?

Many different types of log data are uploaded from Microsoft servers to a proprietary security monitoring solution for near real-time (NRT) analysis and an internal big data computing service (Cosmos) or Azure Data Explorer (Kusto) for long-term storage. This data transfer occurs over a FIPS 140-2-validated TLS connection on approved ports and protocols using automated log management tools.

Logs are processed in NRT using rule-based, statistical, and machine learning methods to detect system performance indicators and potential security events. Machine learning

models use incoming log data and historical log data stored in Cosmos or Kusto to continuously improve detection capabilities. Security-related detections generate alerts, notifying on-call engineers of a potential incident and triggering automated remediation actions when applicable. In addition to automated security monitoring, service teams use analysis tools and dashboards for data correlation, interactive queries, and data analytics. These reports are used to monitor and improve the overall performance of the service.



How do Microsoft online services protect audit logs?

The tools used in Microsoft online services to collect and process audit records don't allow permanent or irreversible changes to the original audit record content or time ordering. Access to Microsoft online service data stored in Cosmos or Kusto is restricted to authorized personnel. In addition, Microsoft restricts the management of audit logs to a limited subset of security team members responsible for audit functionality. Security team personnel don't have standing administrative access to Cosmos or Kusto. Administrative access requires Just-In-Time (JIT) access approval, and all changes to logging mechanisms for Cosmos are recorded and audited. Audit logs are retained long enough to support incident investigations and meet regulatory requirements. The exact period of audit log data retention determined by the service teams; most audit log data is retained for 90 days in Cosmos and 180 days in Kusto.

How do Microsoft online services protect user personal data that may be captured in audit logs?

Prior to uploading log data, an automated log management application uses a scrubbing service to remove any fields that contain customer data, such as tenant information and user personal data, and replace those fields with a hash value. The anonymized and hashed logs are rewritten and then uploaded into Cosmos. All log transfers occur over a TLS encrypted connection (FIPS 140-2).

What is Microsoft's strategy for monitoring security?

Microsoft engages in continuous security monitoring of its systems to detect and respond to threats to Microsoft online services. Our key principles for security monitoring and alerting are:

- Robustness: signals and logic to detect various attack behaviors
- Accuracy: meaningful alerts to avoid distractions from noise
- Speed: ability to catch attackers quickly enough to stop them

Automation, scale, and cloud-based solutions are key pillars of our monitoring and response strategy. For us to effectively prevent attacks at the scale of some of the Microsoft online services, our monitoring systems need to automatically raise highly accurate alerts in near real time. Likewise, when an issue is detected, we need the ability to mitigate the risk at scale, we cannot rely on our team to manually fix issues machine-by-machine. To mitigate risks at scale, we use cloud-based tools to automatically apply countermeasures and provide engineers with tools to apply approved mitigation actions quickly across the environment.

How do Microsoft online services perform security monitoring?

Microsoft online services use centralized logging to collect and analyze log events for activities that might indicate a security incident. Centralized logging tools aggregate logs from all system components, including event logs, application logs, access control logs, and network-based intrusion detection systems. In addition to server logging and application-level data, core infrastructure is equipped with customized security agents

that generate detailed telemetry and provide host-based intrusion detection. We use this telemetry for monitoring and forensics.

The logging and telemetry data we collect enables 24/7 security alerting. Our alerting system analyzes log data as it gets uploaded, producing alerts in near real time. This includes rules-based alerts and more sophisticated alerting based on machine learning models. Our monitoring logic goes beyond generic attack scenarios and incorporates deep awareness of service architecture and operations. We analyze security monitoring data to continuously improve our models to detect new kinds of attacks and improve the accuracy of our security monitoring.

How do Microsoft online services respond to security monitoring alerts?

When security events that trigger alerts require responsive action or further investigation of forensic evidence throughout the service, our cloud-based tools allow for rapid response throughout the environment. These tools include fully automated, intelligent agents that respond to detected threats with security countermeasures. In many cases, these agents deploy automatic countermeasures to mitigate security detections at scale without human intervention. When this response is not possible, the security monitoring system automatically alerts the appropriate on-call engineers, who are equipped with a set of tools that enable them to act in real time to mitigate detected threats at scale. Potential incidents are escalated to the appropriate Microsoft security response team and are resolved using the security incident response process.

How do Microsoft online services monitor system availability?

Microsoft actively monitors its systems for indicators of resource over-utilization and abnormal use. Resource monitoring is complemented by service redundancies to help avoid unexpected downtime and provide customers with reliable access to products and services. Microsoft online service health issues are communicated promptly to customers through the Service Health Dashboard (SHD).

Azure and Dynamics 365 online services utilize multiple infrastructure services to monitor their security and health availability. The implementation of Synthetic Transaction (STX) testing allows Azure and Dynamics services to check the availability of their services. The STX framework is designed to support the automated testing of components in running services and is tested on live site failure alerts. Additionally, the Azure Security Monitoring (ASM) service has implemented centralized synthetic testing

procedures to verify security alerts function as expected in both new and running services.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to audit logging and monitoring.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
ISO 27001/27002 Statement of Applicability Certificate	A.12.1.3: Availability monitoring and capacity planning A.12.4: Logging and monitoring	April 24, 2023
ISO 27017 Statement of Applicability Certificate	A.12.1.3: Availability monitoring and capacity planning A.12.4: Logging and monitoring A.16.1: Management of information security incidents and improvements	April 24, 2023
ISO 27018 Statement of Applicability Certificate	A.12.4: Logging and monitoring	April 24, 2023
SOC 1	IM-1: Incident management framework IM-2: Incident detection configuration IM-3: Incident management procedures IM-4: Incident post-mortem VM-1: Security event logging and collection VM-12: Azure services availability monitoring VM-4: Malicious events investigation VM-6: Security vulnerability monitoring	November 17, 2023
SOC 2 SOC 3	C5-6: Restricted access to logs IM-1: Incident management framework IM-2: Incident detection configuration IM-3: Incident management procedures	November 17, 2023

External audits	Section	Latest report date
	IM-4: Incident post-mortem PI-2: Azure portal SLA performance review VM-1: Security event logging and collection VM-12: Azure services availability monitoring VM-4: Malicious events investigation VM-6: Security vulnerability monitoringVM	

Microsoft 365

Expand table

External audits	Section	Latest report date
FedRAMP (Office 365)	AC-2: Account management AC-17: Remote access AU-2: Audit events AU-3: Content of audit records AU-4: Audit storage capacity AU-5: Response to audit processing failures AU-6: Audit review, analysis, and reporting AU-7: Audit reduction and report generation AU-8: Time stamps AU-9: Protection of audit information AU-10: Non-repudiation AU-11: Audit record retention AU-12: Audit generation SI-4: Information system monitoring SI-7: Software, firmware, and information integrity	July 31, 2023
ISO 27001/27002/27017 Statement of Applicability Certification (27001/27002) Certification (27017)	A.12.3: Availability monitoring and capacity planning A.12.4: Logging and monitoring	March 2023
SOC 1 SOC 2	CA-19: Change monitoring CA-26: Security incident reporting CA-29: On-call engineers CA-30: Availability monitoring CA-48: Datacenter logging CA-60: Audit logging	January 3, 2023

External audits	Section	Latest report date
SOC 3	CUEC-08: Reporting incidents CUEC-10: Service contracts	January 3, 2023

Resources

- Video: Microsoft 365 audit logging and monitoring

Feedback

Was this page helpful? Yes No

Management and operation of the Azure production network

Article • 08/29/2023

This article describes how Microsoft manages and operates the Azure production network to secure the Azure datacenters.

Monitor, log, and report

The management and operation of the Azure production network is a coordinated effort between the operations teams of Azure and Azure SQL Database. The teams use several system and application performance-monitoring tools in the environment. And they use appropriate tools to monitor network devices, servers, services, and application processes.

To ensure the secure execution of services running in the Azure environment, the operations teams implement multiple levels of monitoring, logging, and reporting, including the following actions:

- Primarily, the Microsoft Monitoring Agent (MMA) gathers monitoring and diagnostic log information from many places, including the fabric controller (FC) and the root operating system (OS), and writes it to log files. The agent eventually pushes a digested subset of the information into a pre-configured Azure storage account. In addition, the freestanding monitoring and diagnostic service reads various monitoring and diagnostic log data and summarizes the information. The monitoring and diagnostic service writes the information to an integrated log. Azure uses the custom-built Azure security monitoring, which is an extension to the Azure monitoring system. It has components that observe, analyze, and report on security-pertinent events from various points in the platform.

- The Azure SQL Database Windows Fabric platform provides management, deployment, development, and operational oversight services for Azure SQL Database. The platform offers distributed, multi-step deployment services, health monitoring, automatic repairs, and service version compliance. It provides the following services:
 - Service modeling capabilities with high-fidelity development environment (datacenter clusters are expensive and scarce).
 - One-click deployment and upgrade workflows for service bootstrap and maintenance.
 - Health reporting with automated repair workflows to enable self-healing.

- Real time monitoring, alerting, and debugging facilities across the nodes of a distributed system.
- Centralized collection of operational data and metrics for distributed root cause analysis and service insight.
- Operational tooling for deployment, change management, and monitoring.
- The Azure SQL Database Windows Fabric platform and watchdog scripts run continuously and monitor in real time.

If any anomalies occur, the incident response process followed by the Azure incident triage team is activated. The appropriate Azure support personnel are notified to respond to the incident. Issue tracking and resolution are documented and managed in a centralized ticketing system. System uptime metrics are available under the non-disclosure agreement (NDA) and upon request.

Corporate network and multi-factor access to production

The corporate network user base includes Azure support personnel. The corporate network supports internal corporate functions and includes access to internal applications that are used for Azure customer support. The corporate network is both logically and physically separated from the Azure production network. Azure personnel access the corporate network by using Azure workstations and laptops. All users must have an Azure Active Directory (Azure AD) account, including a username and password, to access corporate network resources. Corporate network access uses Azure AD accounts, which are issued to all Microsoft personnel, contractors, and vendors and managed by Microsoft Information Technology. Unique user identifiers distinguish personnel based on their employment status at Microsoft.

Access to internal Azure applications is controlled through authentication with Active Directory Federation Services (AD FS). AD FS is a service hosted by Microsoft Information Technology that provides authentication of corporate network users through applying a secure token and user claims. AD FS enables internal Azure applications to authenticate users against the Microsoft corporate active directory domain. To access the production network from the corporate network environment, users must authenticate by using multi-factor authentication.

Next steps

To learn more about what Microsoft does to secure the Azure infrastructure, see:

- Azure facilities, premises, and physical security
- Azure infrastructure availability
- Azure information system components and boundaries
- Azure network architecture
- Azure production network
- Azure SQL Database security features
- Azure infrastructure monitoring
- Azure infrastructure integrity
- Azure customer data protection

Microsoft 365 audit log collection

Article • 03/02/2023

Audit logs serve an important role in maintaining, troubleshooting, and protecting both customer tenants and the internal Microsoft 365 infrastructure. Due to the scale at which Microsoft 365 operates, the collection and processing of audit logs must be strategically managed to ensure efficient and effective monitoring. This article provides an overview of Microsoft 365's audit logging practices including what event types are captured, what information is contained in each log, how logging policies are enforced, and how log data is protected at all stages of its lifecycle.

Defining auditable events

The Microsoft 365 Security team is responsible for defining the baseline logs that must be collected across Microsoft 365. They maintain an official list of event types that each system must capture as well as the data each logged event must contain.

Microsoft 365 captures log data from various sources, including:

- Event logs
- AppLocker logs
- Performance data
- System Center data
- Call detail records
- Quality of experience data
- IIS Web Server logs
- SQL Server logs
- Syslog data
- Security audit logs

Each of the logs in this list must at least contain the following information to establish event type, source, location, outcome, time, and entity associated:

- Source User ID
- Target User ID – as relevant/appropriate for event type
- Event timestamp (Date & Time)
- Event details
 - Type
 - Outcome (success/failure)
 - Detail information – defined per event type
- Source & Target Hostname – as relevant/appropriate for event type

- Source & Target network addresses and protocols – as relevant/appropriate for event type

The list of auditable events and related data is informed by ongoing risk assessments, Microsoft 365 security standards, business requirements, and compliance requirements. The Microsoft 365 Security team reviews and updates the list of auditable events to account for new threats, system changes, lessons learned from past incidents, and changing compliance requirements.

Service teams can define additional logging requirements for their service in addition to the list of auditable events defined by the Microsoft 365 Security team. Application-specific events are reviewed and updated during service reviews and the planning phases of feature milestones. The Microsoft 365 Security team also helps guide these individual service teams on audit functions to meet their specific needs. Service-level auditable events are reviewed and updated whenever a significant change to the system is made.

Due to Microsoft's scale, the amount of data captured must be balanced with the ability to store and process it. Collecting information on every possible event and the contents therein would be impractical from both a resource and analytical standpoint. By being selective with the kinds of log data collected, Microsoft can maintain the health and security of its information systems in an efficient and effective manner.

Centralized management

Microsoft 365 enforces logging requirements centrally through the policies set by the Microsoft 365 Security Team. Each Microsoft 365 system must include a custom logging agent, Office Data Loader (ODL), which is installed as part of the system baseline. ODL enforces central logging policies and is configured to collect and send the events defined by Microsoft 365 Security to centralized services for processing and storage.

ODL is configured to scrub all end-user information automatically and upload events in batches every five minutes. Any fields that contain customer data, such as tenant information and end-user identifiable information, are removed and replaced with a hash value. Permanent or irreversible changes to audit record content and time ordering, aside from the scrubbing described previously, is prohibited.

Log data is uploaded to a proprietary security monitoring solution for near real-time (NRT) analysis, checking logs for potential security events and performance indicators. Logs are also uploaded to an internal, big data computing service (Azure Data Lake) for long-term storage. All log data transfers occur over a FIPS 140-2-validated TLS connection to ensure data confidentiality during transit.

Most audit log data stored in Azure Data Lake are kept for at least one year to support investigations of security incidents and to meet regulatory retention requirements. Service teams may select alternative retention periods of 90 days or longer for specific types of log data to support the needs of their applications. Microsoft 365 log retention and backup policies ensure log data is readily available for incident investigations, compliance reporting, and any other business requirements.

Access control

Microsoft performs extensive monitoring and auditing of all delegation, privileges, and operations that occur within Microsoft 365. Access to Microsoft 365 data stored in Azure Data Lake is restricted to authorized personnel, and all access control requests and approvals are captured for security event analysis. Microsoft reviews access levels in near real-time to ensure that its systems can only be accessed by users who have authorized business justifications and meet the eligibility requirements. All permitted access is traceable to a unique user.

Microsoft restricts the management of audit logs to a limited subset of Security Team members responsible for audit functionality. Microsoft's internal access control philosophy revolves around the principles of Just-In-Time (JIT) and Just-Enough-Access (JEA). As such, the Security Team does not have standing administrative access to Azure Data Lake, and all changes are recorded and audited.

Auditing and reporting in Microsoft cloud services

Article • 04/04/2023

Microsoft cloud services include several auditing and reporting features you can use to track user and administrative activity within their tenant, Examples include changes made to Exchange Online and SharePoint Online tenant configuration settings, and changes made by users to documents and other items. You can use audit information and reports available in Microsoft cloud services to more effectively manage user experience, mitigate risks, and fulfill compliance obligations.

Security & Compliance Centers

The Microsoft 365 Security & Compliance Center, Microsoft Defender portal, and Microsoft Purview compliance portal are one-stop portals for protecting data in your organization, and they include many auditing and reporting features. These centers help you with your data protection or compliance needs and audit user and administrator activity. You can access these centers using your subscription admin account.

These centers include navigation panes for access to several features:

- **Alerts:** Enables you to manage alerts, view security-related alerts, and manage advanced alerts using Defender for Cloud Apps.
- **Permissions:** Enables you to assign permissions such as Compliance Administrator, eDiscovery Manager, and others to people in your organization so they can perform tasks in these centers. You assign permissions for most features in each center, but other permissions must be configured using the Exchange admin center and SharePoint admin center.
- **Threat management:** Enables you to create and apply device management policies using Basic Mobility and Security for Microsoft 365, to set up data loss prevention (DLP) policies for your organization, to configure email filtering, anti-malware, DomainKeys Identified Mail (DKIM), safe attachments, safe links, and OAuth apps.
- **Data governance:** Enables you to import email or SharePoint data from other systems into Microsoft 365, configure archive mailboxes, and set retention policies for email and other content within your organization.
- **Search & investigation:** Provides content search, audit log, quarantine, and eDiscovery case management tools to quickly drill into activity across Exchange

Online mailboxes, groups and public folders, SharePoint Online, and OneDrive for Business.

- **Reports:** Enables you to quickly access reports for SharePoint Online, OneDrive for Business, Exchange Online, and Microsoft Entra ID.
- **Service assurance:** Provides information about how Microsoft maintains security, privacy, and compliance with global standards for Microsoft 365, Azure, Microsoft Dynamics CRM Online, Microsoft Intune, and other cloud services. Also includes access to third-party ISO, SOC, and other audit reports, as well as Audited Controls, which provides details about the various controls that have been tested and verified by third-party auditors of Microsoft 365.

Service assurance

Many organizations in regulated industries are subject to extensive compliance requirements. To perform their own risk assessments, customers often need in-depth information about how Microsoft 365 maintains the security and privacy of their data. Microsoft is committed to the security and privacy of customer data in its cloud services and to earning customer trust by providing a transparent view of its operations, and easy access to independent compliance reports and assessments.

Service Assurance provides transparency of operations and information about how Microsoft maintains the security, privacy, and compliance of customer data in Microsoft 365. It includes third-party audit reports along with a library of white papers, FAQs, and other materials on Microsoft 365 topics such as data encryption, data resiliency, security incident management and more. Customers can use this information to perform their own regulatory risk assessments. Compliance officers can assign the "Service Assurance User" role to give users access to Service Assurance. The tenant administrator can also provide external users, such as independent auditors, with access to information in the Service Assurance dashboard through the Microsoft Cloud Service Trust Portal (STP).

Resources

- Microsoft 365 Reporting Features
- Internal Logging for Microsoft 365 Engineering

Microsoft 365 reporting features

Article • 10/20/2023

Reporting features in Microsoft 365 provides various audit reports for Microsoft Entra ID, Exchange Online, device management, supervisory review, and data loss prevention (DLP). These reports are different and separate from the Microsoft 365 activity reports.

Microsoft 365 Reports dashboard

The Reports dashboard in the Microsoft 365 admin center preview displays usage activity across Microsoft 365. Microsoft 365 global administrators, or an Exchange Online, SharePoint Online, or Skype for Business administrator, can get granular insight into the usage of that service. For example, the number of users in a particular Microsoft 365 service, the number of users that have activated Microsoft 365 Apps for enterprise (previously named Office 365 ProPlus), and how much mail is flowing through the organization. Reports are available for the last 7, 30, 90, and 180 days.

The following reports are available:

- Email activity report
- Microsoft Office activations report
- SharePoint Online Site usage report
- OneDrive for Business usage report
- Viva Engage activity report
- Skype for Business activity report
- Skype for Business Peer-to-Peer activity report
- Skype for Business Conference Organizer report
- Skype for Business Conference Participant activity report

For more information, see Activity Reports in the Microsoft 365 admin center.

Microsoft Entra reports

Microsoft 365 uses Microsoft Entra ID for authentication and identity management. Microsoft 365 administrators use reports generated by Azure to identify unusual activity and unauthorized access to their data. You can use access and usage reports in Microsoft Entra ID to gain visibility into the directory integrity and security for your organization. With this information, you can identify and mitigate possible security risks.

Microsoft Entra reports can be exported to Microsoft Excel and correlated with other data from Microsoft 365. For example, the results of an audit log search can provide insight into access, authentication, and application-level activities. Advanced anomaly and resource usage reports are available with Microsoft Entra ID P1 or P2. These advanced reports help improve your security posture and help you respond to potential threats by applying analytics about device access and application usage. For more information, see Microsoft Entra ID reporting.

Exchange Online Audit Reports

Exchange Online audit reports include details on mailbox access and changes made by administrators to your Exchange Online tenant. With mailbox auditing, you can use the tasks in the following table to run reports and export Exchange Online audit logs.

> **Note**
>
> You must enable mailbox audit logging for each mailbox so that audited events are saved in the audit log for that mailbox. If mailbox audit logging isn't enabled for a mailbox, events for that mailbox won't be saved in the audit log and won't appear in mailbox audit reports. For more information, see **enable mailbox auditing**.

Task	Description
Run a non-owner mailbox access report	Displays the list of mailboxes accessed by someone other than the owner of the mailbox. The report contains information about who accessed the mailbox, the actions they took in the mailbox, and whether the actions were successful.
Export mailbox audit logs	Mailbox audit logs contain information on access and actions in a mailbox taken by a user other than the mailbox owner. Administrators can specify mailboxes along with a date range to generate reports. The logs are exported in XML, attached to a message, and sent to specific users as determined by the administrator.
Run an administrator role group report	The administrator role group assigns administrative privileges to users. These privileges allow users to perform administrative tasks such as reset passwords, create or modify mailboxes, and assign admin privileges to other users. The admin role group report shows changes to role groups, including the addition or removal of members.
View the admin audit log	The admin audit log report lists all create, update, and delete functions performed by administrators in Exchange Online. Log entries provide information on which cmdlet was run, what parameters were used, who ran the cmdlet, and what objects were affected.

Task	Description
Mailbox content search and hold	Provides details of any changes to In-Place eDiscovery or In-Place Hold settings on mailboxes.
Export the admin audit log	The admin audit log records specific administrative actions such as create, update, and delete in Exchange Online. The results from the log are exported to XML and administrators can choose to send this log to a set of users.
Run a per-mailbox litigation hold report	Provides details of any changes to litigation hold settings on mailboxes.
View and export the external admin audit log	Contains details of actions performed by external administrators. The entries provide information on which cmdlet was run, what parameters were used, and any actions that create, modify, or delete objects in Exchange Online.

Device compliance Reports

You manage and secure mobile devices connected to your subscription using Basic Mobility and Security for Microsoft 365. Mobile devices used to access work email, calendar, contacts, and documents play a significant part in making sure that employees are able to work anytime and from anywhere. It's critical that you protect your organization's information. You use Basic Mobility and Security for Microsoft 365 to set device security policies and access rules. If lost or stolen, you also use Basic Mobility and Security for Microsoft 365 to wipe mobile devices.

Basic Mobility and Security compliance reports provide an overview of policies set up by an organization to secure mobile devices that access Microsoft 365 data. The report allows filtering of devices by compliance status, reported violations, blocked devices, and how many devices wiped as a result of security policies. For more information, see Overview of Basic Mobility and Security for Microsoft 365 .

Data loss prevention

DLP policies help manage the security and flow of information in an organization. You can set up policies to block access to content, encrypt data, or notify users of policy and policy violations using in-application DLP Policy Tips. DLP reports provide insight into the number of policy and rule matches, overrides, and false positives.

You use the Microsoft 365 admin center to view information about the number of messages detected by your DLP policies. DLP reports provide insights into policy and

rule matches for sent and received mail. You can also view the number of matches, overrides, and false positives for each policy within the past 24 hours using the Exchange admin center. If you download an Excel report, you can view even more detail, such as who sent which message, on what day, and what policy matches were triggered. For more information, see View reports about DLP policy detections.

Auditing in Viva Engage

Viva Engage provides administrators with the ability to export user activity data from their Viva Engage network via the Viva Engage data export API, or manually via the Viva Engage network admin page. The ability to export logs is restricted to Network Administrators in Viva Engage. (All Microsoft 365 global administrators are Viva Engage Network Administrators.)

The following data is exportable:

Filename	Description
Users.csv	All new, pending, and suspended users in the network
Messages.csv	All messages in the network
Files.csv (metadata)	Metadata such as filename, file API URL, uploader ID, uploaded at, etc.
Files.csv (Original files)	Zip file of the original files uploaded by users into Viva Engage
Topics.csv	Topics created on the network
Pages.csv	Pages (notes) created by users in the network
Admins.csv	All verified administrators on the network
Networks.csv	All Viva Engage external networks

Viva Engage data is also available through the Microsoft 365 activity reports. In addition, Viva Engage is actively working on exposing additional logging via the Microsoft 365 Management Activity API, and on the ability to reason over data using Power BI. See the Office Roadmap for more information on these features.

Datacenter security overview

Article • 05/15/2023

How does Microsoft host its online services?

Microsoft delivers more than 200 cloud services, including enterprise services such as Microsoft Azure, Microsoft 365, and Microsoft Dynamics 365, to customers 24x7x365. These services are hosted in Microsoft's cloud infrastructure composed of globally distributed datacenters, edge computing nodes, and service operations centers. They are supported and connected by one of the world's largest global networks, with an extensive fiber footprint.

The datacenters that power our cloud offerings focus on high reliability, operational excellence, cost-effectiveness, environmental sustainability, and a trustworthy online experience for customers and partners worldwide. Microsoft regularly tests our datacenter security through both internal and third-party audits. As a result, the most highly regulated organizations in the world trust the Microsoft cloud, which is compliant with more certifications than any other cloud service provider.

How does Microsoft protect its datacenters from unauthorized access?

Access to physical datacenter facilities is tightly controlled by outer and inner perimeters with increasing security at each level, including perimeter fencing, security officers, locked server racks, integrated alarm systems, around-the-clock video surveillance by the operations center, and multi-factor access control. Only required personnel are authorized to access Microsoft datacenters. Logical access to Microsoft 365 infrastructure, including customer data, is prohibited from within Microsoft datacenters.

Our Security Operations Centers use video surveillance along with integrated electronic access control systems to monitor datacenter sites and facilities. Cameras are strategically positioned for effective coverage of the facility perimeter, entrances, shipping bays, server cages, interior aisles, and other sensitive security points of interest. As part of our multi-layered security posture, any unauthorized entry attempts detected by the integrated security systems generate alerts to security personnel for immediate response and remediation.

How does Microsoft protect its datacenters from environmental hazards?

Microsoft employs a variety of safeguards to protect against environmental threats to datacenter availability. Datacenter sites are strategically selected to minimize risk from a variety of factors, including floods, earthquakes, hurricanes, and other natural disasters. Our datacenters use climate control to monitor and maintain optimized conditioned spaces for staff, equipment, and hardware. Fire detection and suppression systems and water sensors help to detect and prevent fire and water damage to equipment.

Disasters are unpredictable, but Microsoft datacenters and operations personnel prepare for disasters to provide continuity of operations should unexpected events occur. Resilient architecture and up-to-date tested continuity plans mitigate potential damage and promote swift recovery of datacenter operations. Crisis management plans provide clarity on roles, responsibilities, and mitigation activities before, during, and after a crisis. The roles and contacts defined in these plans facilitate effective escalation up the chain of command during crisis situations.

How does Microsoft verify the effectiveness of datacenter security?

We understand that for our customers to fully realize the benefits of the cloud, they must be able to trust their cloud service provider. Our infrastructure and suite of cloud services are built from the ground up to address the rigorous security and privacy requirements of our customers. We help our customers comply with national, regional, and industry-specific requirements governing the collection and use of individuals' data by providing the most comprehensive set of compliance offerings of any cloud service provider.

Our cloud infrastructure and offerings meet a broad set of international and industry-specific compliance standards, such as ISO, HIPAA, FedRAMP, and SOC, as well as country-specific standards, like Australia's IRAP, UK's G-Cloud, and Singapore's MTCS. Rigorous, third-party audits verify our adherence to the strict security controls these standards mandate. Audit reports for our datacenter infrastructure and cloud offerings are available at the Microsoft Service Trust Portal.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to

datacenter security.

Expand table

External audits	Section	Latest report date
ISO 27001/27002 (Azure) Statement of Applicability Certificate	A.11: Physical and environmental security	April 24, 2023
SOC 1 (Azure)	PE-1: Datacenter physical access provisioning PE-2: Datacenter security verification PE-3: Datacenter user access review PE-4: Datacenter physical access mechanisms PE-5: Datacenter physical surveillance monitoring PE-6: Datacenter critical environment maintenance PE-7: Datacenter environmental controls PE-8: Datacenter incident response	November 17, 2023
SOC 2 (Azure)	PE-1: Datacenter physical access provisioning PE-2: Datacenter security verification PE-3: Datacenter user access review PE-4: Datacenter physical access mechanisms PE-5: Datacenter physical surveillance monitoring PE-6: Datacenter critical environment maintenance PE-7: Datacenter environmental controls PE-8: Datacenter incident response	November 17, 2023

Resources

- Take a virtual tour of a Microsoft datacenter
- Interactive: Global Azure Infrastructure

Feedback

Was this page helpful? Yes No

Datacenter architecture and infrastructure

Article • 03/02/2023

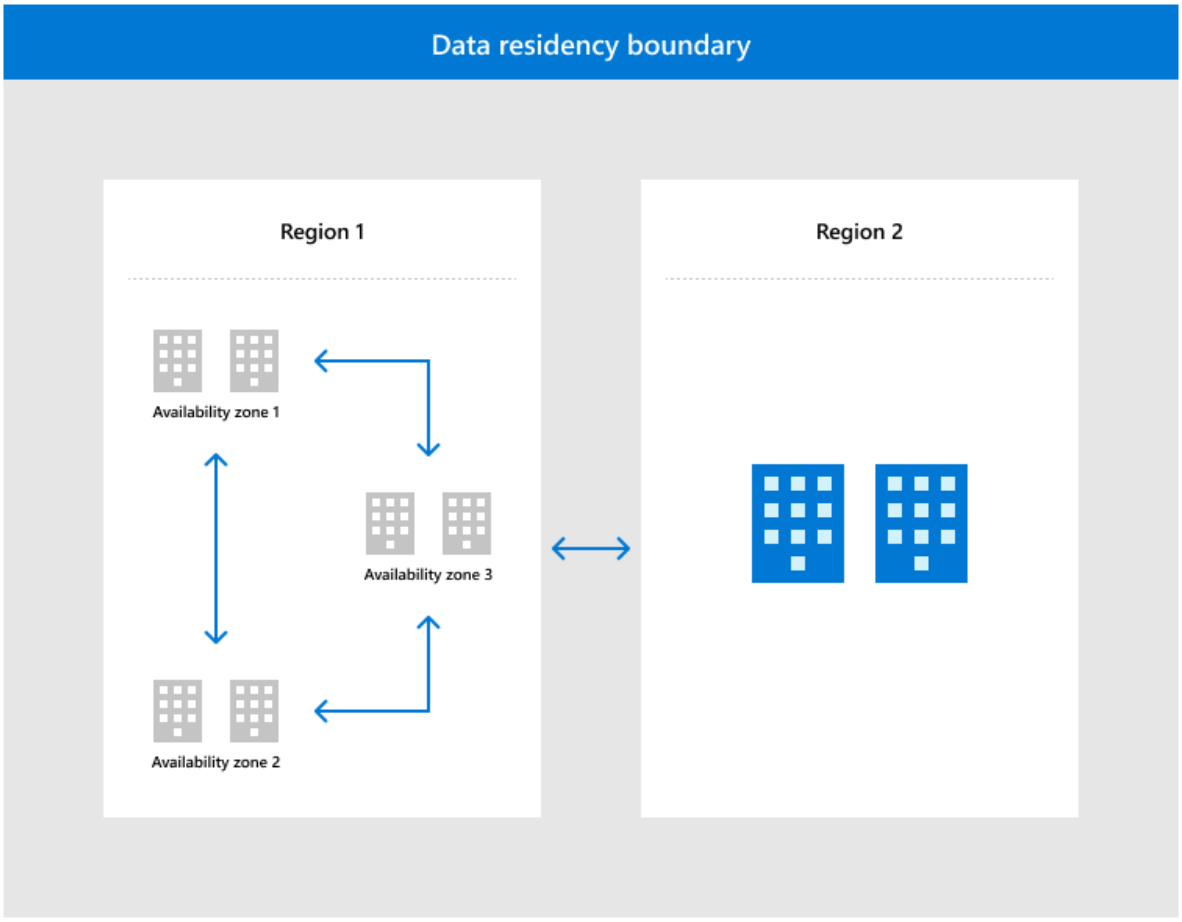
Microsoft datacenters are designed to implement a strategy of defense-in-depth, employing multiple layers of safeguards to reliably protect our cloud architecture, and supporting infrastructure. Redundancy is built into all systems at multiple levels to support datacenter availability.

Microsoft has highly secured datacenter facilities spread worldwide creating a distributed datacenter infrastructure, supporting thousands of online services. This globally distributed infrastructure is designed to bring applications closer to users, preserve data residency, and offer comprehensive compliance and resiliency options for customers.

Regions are sets of datacenters that are interconnected via a massive and resilient network. Regions are organized into geographies, granting customers with specific data-residency and compliance needs the ability to keep their data and applications close. Built-in fault tolerance allows geographies to withstand complete region failure through their connection to the dedicated, high-capacity networking infrastructure.

Physically separate locations within a region are referred to as availability zones, each being made up of one or more datacenters equipped with independent power, cooling, and networking. Availability zones allow for mission-critical applications to run with high availability and low-latency replication.

The following figure shows how the global infrastructure pairs region and availability zones within the same data residency boundary for high availability, disaster recovery, and backup.



Geographically distributed datacenters enable Microsoft to bring services closer to customers, reduce network latency, and allow for geo-redundant backup and failover.

Availability

Microsoft datacenters are engineered to provide 99.999% availability to meet customer SLAs and service needs. Microsoft invests significantly in the global operations, management, networks, and sustainability of facilities that deliver services 24x7x365.

Compliance Standards and Requirements

Microsoft has invested over $15 billion building our global infrastructure and over $9 billion in research and development to increase efficiency and drive innovation. As a result, Microsoft's datacenters are evolving at a more rapid pace than many facilities in the industry, and thus do not follow prescriptive requirements outlined by traditional datacenter standards. In addition to the wealth of operational insight that comes with running one of the world's largest datacenter portfolios, Microsoft uses IEEE Gold Book data and third-party reliability simulation software to continuously improve our datacenter design standards. Microsoft datacenters are audited extensively as part of several regulatory audits as called out in the compliance portfolio. The level of maturity

in Microsoft datacenters can be evaluated through the compliance portfolio and, for resiliency specifically, the ISO 22301 certification.

While Microsoft operates programs in alignment with the spirit of the ANSI/TIA-942 Telecommunications Infrastructure of Datacenters Standard, portions of this standard are not applicable to Microsoft or conflict with other regulatory and/or country-specific requirements. Moreover, Microsoft has chosen to use a more performance-based approach to match customer needs.

Data and network redundancy

Critical datacenter facilities employ multiple layers of redundant systems to sustain faults and minimize service disruptions. Locally redundant storage at the disk level protects data within a region, with geo-redundant storage providing intra-region redundancy. To ensure reliable network communications, Microsoft owns and utilizes diverse fiber routes and redundant hardware to protect critical components from failure or service disruption.

Geo-replication is used to provide redundancy to alternate geographic locations. Data durability is obtained by synchronously replicating data across multiple databases in different datacenters. Restoration tests are performed for all backup data owned by the cloud. Disaster Recovery is achieved by asynchronous replication to a datacenter in a different geographical region.

Capacity

Cloud Operations is a dedicated capacity team that forecasts future requirements to ensure necessary capacity is structured and available for customer and internal use. Systems are monitored to ensure acceptable service performance, availability, service utilization, storage utilization, network latency, and audit log capacity. Microsoft also protects datacenters against the effects of denial-of-service attacks on bandwidth, transactional capacity, and storage capacity.

All service teams include capacity planning as a key feature of their datacenter models and data replication plans to ensure that there is necessary capacity for information processing, telecommunications, and environmental support.

Power

Microsoft's datacenters have dedicated 24x7 uninterruptible power supplies (UPSs) and emergency power support, which includes on-site generators that provide backup

power. Regular maintenance and testing are conducted for both the UPS and generators, and operations teams have contractual agreements with local vendors for emergency fuel delivery. Datacenters also have a dedicated Facility Operations Center to monitor power systems, including critical electrical components.

Microsoft datacenters are equipped with protective spaces and appropriate labeling for cables. The power infrastructure equipment is placed in environments that have been engineered to protect from environmental risks. All portable online services' assets must be locked or fastened in place in order to provide protection against theft or movement damage. Power cables are run under the floors, overhead in cable trays, and within cabinets for protection from moving parts and accidental damage. All electrical spaces are behind card readers or additional key locks as appropriate. Access hallways, exterior entrances, and equipment yards are all monitored via video surveillance. Power systems also utilize redundancy as a form of protection, with multiple power/utility feeds into the facility and generators and UPS systems.

A long-term alternate power supply is implemented for the information system that is capable of maintaining power in a minimally required operational capability. When power fails or drops to an unacceptable voltage level, UPS systems instantly come online. This provides enough power for running the servers until the generators can take over. Emergency generators provide back-up power for extended outages, planned maintenance and can operate the datacenter with on-site fuel reserves if a natural disaster occurs.

Microsoft datacenters (both leased and fully managed) implement emergency lighting in the form of overhead emergency lighting on dedicated circuits backed up by UPS and generator systems. Automatic emergency lighting is implemented in accordance with the National Fire and Protection Association (NFPA) Life Safety Code or the applicable local code/law. If utility power is lost, the emergency lighting will automatically switch to power provided by the UPS and generator systems. The emergency lighting systems within the datacenters undergo routine maintenance to ensure that they remain in proper working order.

Maintenance

System maintenance policy and procedures are in place in accordance with *Microsoft's Online Services Physical and Environmental Security Standard*. All Microsoft equipment and systems are regularly maintained in order to guarantee operational efficiency. Maintenance of any equipment or system must be performed in accordance with manufacturer's recommendations, carried out by authorized personnel, and recorded in a maintenance ticket.

There are two asset teams that maintain different types of systems:

- **Critical Environment (CE)** team:
 - CE is the team that provides facility management for electrical, mechanical, and physical systems that comprise the operating infrastructure of the facility. The CE team schedules, performs, documents, and reviews all maintenance activities performed on CE components. Microsoft datacenters rely on a computerized system to manage maintenance schedules and work orders.
 - Datacenter management (DCM) is responsible for all CE maintenance that is performed either on site or remotely. CE maintenance is prescribed in required step-by-step documents called Methods of Procedure (MOP). MOPs are reviewed/approved by datacenter management prior to any work beginning.

- **Site Services** team:
 - Site Services is the team that provides the servicing of Microsoft online service assets located at the Microsoft datacenter. The DC Site Services team provides a smart hands/break fix service for assets belonging to properties provisioning services from the datacenter. For example, assets requiring physical maintenance could request smart hands service from the DC Site Services team. All Site Services work on Microsoft assets is scheduled, performed, documented, and reviewed in work tickets within the workflow ticketing tool, and no work can occur without an approved work ticket.
 - The Technical Program Manager (TPM) and DCM team are responsible for all Site Services work that occurs in the datacenter and work that requires the asset to be transferred offsite. Site Services maintenance is performed in areas of the datacenter that are controlled and protected by physical security mechanisms.

If CE components are required to be removed from the facility, handling the equipment is approved by DCM. In most instances, CE components receive onsite maintenance and are not removed from the facility. Property assets (for example, network devices or servers) requiring transfer offsite must have explicit asset owner approval.

Digital media within the cloud may not be transported from the colocation space unless it is moved to be destroyed. When these assets are to be destroyed, they are stored in locked storage bins that are under CCTV camera coverage. When the assets are ready to be destroyed, a physical security officer and a Microsoft full-time employee from Asset Management must escort the locked bin from the colocation space to where the onsite shredding is to occur. As shredding occurs at the datacenter and under Microsoft supervision, Microsoft assets do not leave the controlled areas of the datacenter.

All maintenance work must be approved prior to work beginning, including access to system maintenance tools. Microsoft Infrastructure has implemented the maintenance

tools control by creating an access level within the Datacenter Access Tool (DCAT). Each facility contains a restricted physical lock box or access-controlled room for the storage of specialized maintenance tools. Access to the lock box or storage room is controlled in the DCAT tool to prohibit unauthorized access to the maintenance tools. This program ensures that only personnel with approved access can access the tools. The Site Services team performs routine inventory checks to verify the status of all tools. On a quarterly basis, the datacenter management team and physical security teams perform audits of the DCAT access list to keep the access list of maintenance personnel current. Personnel terminations or transfers are reflected immediately through a manual update of the access list. Access to the lock box or maintenance storage room is tracked in the access badge reader logs, which are available for any investigations.

The Site Services team maintains an inventory of approved maintenance tools for use within the datacenter. Maintenance personnel are directed to use the provided maintenance tools. Datacenter Management (DCM) approval is required in order to use tools not provided by the datacenter. Physical hand tools are exempt from this type of control.

Microsoft datacenters maintain resident maintenance personnel to support critical datacenter infrastructure systems (the Critical Environment team) and datacenter operations (the Site Services team). The Critical Environment and Site Services teams have identified critical security and technology system components that they maintain spares for onsite. Critical information system services are provisioned from more than one datacenter to prevent an interruption in service due to an incident at one of the datacenters.

Datacenter threat, vulnerability, and risk assessment

Article • 04/04/2023

Microsoft delivers more than 200 cloud services to customers 24x7x365. Some examples are enterprise services such as Microsoft Azure, Microsoft 365, Microsoft Dynamics, and consumer services such as Bing, MSN, Outlook.com, Skype, and Xbox Live. These services are hosted in Microsoft's cloud infrastructure: globally distributed datacenters, edge computing nodes, and service operations centers, and one of the world's largest global networks; with an extensive fiber footprint connecting them all. Since opening the first datacenter in 1989, Microsoft has invested billions of dollars in our infrastructure and remains focused on delivering reliable, scalable, and security-enhanced online services, while efficiently managing operations and costs as services grow.

Microsoft cloud services are built on a foundation of trust and security, with a priority focus on protecting customer data and applications in the cloud with state-of-the-art technology, processes, and encryption. Customer data is stored in Microsoft datacenters that are geographically distributed and protected by layers of in-depth logical and physical security measures. Microsoft datacenters are designed and operated to protect services and data from harm by natural disasters, environmental threats, or unauthorized access.

Threat, vulnerability, and risk assessment methodology

The Threat, Vulnerability, and Risk Assessment (TVRA) program helps you understand how Microsoft identifies and mitigates the impact of physical and environmental threats to Microsoft datacenters. Microsoft is committed to continually updating its risk assessments and methodologies for improvements and as conditions change. As a result, TVRA analyses and conclusions are subject to change and the reports are considered point-in-time.

Microsoft facilitates the TVRA process following these steps:



Risk identification

TVRAs consider a wide range of threat scenarios arising from natural and human-created (including accidental) hazards. The results will vary depending on datacenter location, design, scope of services, and other factors. The TVRA selects the threat scenarios to highlight in the TVRA document based on customer requirements, an independent country, city, and site-level assessment of the risk environment provided by 3rd-party and 1st-party risk information. For regions that have multiple datacenters, TVRA ratings are aggregated to ensure a holistic view of the physical and environmental threats, vulnerabilities, and risks for the locations being assessed.

Types of threat scenarios assessed for datacenter TVRAs include:

- **External threats**: incidents resulting from external intentional or accidental human activities. For example, civil disorder, terrorism, criminal activity, external theft, improvised explosive devices, armed attacks, arson, unauthorized entry, and airplane crashes.
- **Internal threats**: incidents resulting from internal intentional or accidental human activities. For example, internal theft and sabotage.
- **Natural hazards**: a natural process or phenomenon that could negatively impact datacenters. For example, tropical storms, cyclones, floods, landslides, drought, wildfire, earthquakes, volcanic activity, and severe storms with lightning, hail, strong winds, or heavy rain.
- **Environmental threats**: environmental conditions that could negatively impact datacenters. For example, water stress, heat stress, and pandemics.

Risk analysis

Threats are evaluated based on an assessment of their inherent risk; inherent risk is calculated as a function of inherent impact of a threat and inherent likelihood of the threat occurrence in the absence of management action and controls. These

assessments are informed by both internal subject-matter expert (SME) feedback and using external risk indices.

Residual risk

Residual risk is determined as a measure of remaining risk levels after consideration of control effectiveness. Control effectiveness is evaluated as a measure of current management actions and controls designed to prevent or detect threats while assessing the likelihood that the controls will have their desired effect as designed and implemented. These assessments are informed by an aggregation of internal SME feedback on control effectiveness for the datacenter locations addressed in the TVRA.

Report

Once the assessment is completed, a TVRA report is generated for management approval and to support our overall efforts related to risk management.

Resources

- Azure Global Infrastructure

Datacenter environmental safeguards

Article • 03/02/2023

Microsoft employs various safeguards to protect against environmental threats to datacenter availability. These safeguards enable Microsoft to provide secure and available cloud platforms that are trusted by customers.

Site selection

Protecting Microsoft datacenters from environmental hazards begins with site selection. Microsoft datacenter locations undergo stringent assessment prior to selection for datacenter construction. Microsoft datacenter sites are strategically selected to minimize risk from several factors, including floods, earthquakes, hurricanes, and other natural disasters. In addition to minimizing environmental risk, access to low-cost energy and reliable telecommunications infrastructure are among the key factors that influence datacenter site selection.

Climate control

Climate control is a fundamental component of the critical infrastructure within a datacenter and is utilized to monitor and maintain optimized conditioned spaces for both staff and equipment/hardware. Heat-load (as a biproduct of energy consumption) and humidity need to be managed to ensure suitable operational conditions through mechanical intervention. Local climate conditions, various regulatory requirements, and constraints will dictate the most efficient way for it to be achieved.

Temperature and humidity levels are maintained in accordance with the environmental requirements of the IT hardware that is expected in each datacenter. Microsoft datacenters maintain an operating level agreement with their customers such that optimal efficiency is met while maintaining minimum environmental requirements. The temperature and humidity levels are monitored continuously by the datacenter's Building Management System (BMS). Critical Environments (CE) team members monitor the BMS from the Facilities Operations Center (FOC), so that they can manage the temperature and humidity within the datacenter before any alarm points are exceeded. BMS is configured to notify the CE team when certain markers are reached, who then investigate and make adjustments to remediate the climate issue. Acceptable ranges for temperature and humidity are consistent with American Society of Heating, Refrigerating, and Air-conditioning Engineers (ASHRAE) guidelines or similar locally applicable guidelines. Datacenter humidity is measured by Relative Humidity

percentage, Non-Condensing with the current range between 40%-55%. Temperature range is typically between 18 degrees Celsius and 27 degrees Celsius (between 64.4 degrees Fahrenheit and 80.6 degrees Fahrenheit).

Fire and water damage protection

Microsoft employs state-of-the-art fire detection and suppression systems within each datacenter facility. Fire prevention systems are supported by an independent energy source to ensure protection of Microsoft employees and infrastructure if there is a fire. Datacenter facilities are also protected from water damage by water sensors placed in areas deemed to have leakage risk. These water sensors quickly notify appropriate personnel if there is a water-related emergency. Water shutoff valves are designed to be accessible, and employees are trained in their operation and location.

Microsoft datacenters implement robust fire detection mechanisms including photoelectric smoke detectors installed below the floor and on the ceiling, Xtralis Very Early Smoke Detection Apparatus (VESDA) systems in each colocation, pull station fire alarm boxes installed throughout the datacenters, fire extinguishers located throughout the datacenters, security staff patrols in all building areas multiple times every eight-hour shift, and fire detection/suppression and emergency lighting systems are wired into the datacenter backup power systems providing a redundant power source. Areas containing sensitive electrical equipment are protected by double interlock pre-action (dry pipe) sprinkler systems.

The CE team does a daily site walk-through (DSWT) to check each room and many component parts within them to ensure all fire watch requirements are being met.

Microsoft employs fire detection devices/systems for the information system that activate automatically and notify datacenter personnel along with emergency responders if a fire occurs. If one of the fire detection mechanisms is activated in any colocation space, the local fire department and the Global Security Operations Center in Redmond, Washington are automatically notified. Fire protection and detection systems are tied into the security system notifying the local facility and security staff.

Microsoft provides water/leak detection in areas with a risk of water leakage. Fire suppression systems also have leak detection alarms that are monitored. The water/leak detection system is integrated with the facility alarm and notification system, and sprinkler systems within the datacenters are zoned. The CE and Datacenter Management teams are familiar with the emergency procedures requiring the use of the water shutoff valves and their locations. Sprinkler risers can be shut off individually or as a group via gate valves. All sprinklers in the critical space are double interlock pre-action type

sprinklers that require two forms of activation before flow is initiated. The pressure of the sprinkler system is monitored and alarmed against water leakage.

Health and safety

The commitment to the health and safety of Microsoft employees is fundamental to how success is measured. Microsoft datacenters operate in accordance with local, regional, national, and federal health and safety regulations. In most cases, Microsoft health and safety policies exceed government and regulatory requirements to ensure protection of the health, safety, and wellbeing of all datacenter workers during all phases of the datacenter lifecycle, from construction to decommissioning. Safety management plans, risk assessment, high-risk activity management, and safety training are central tenets to providing a safe work environment for all employees.

Energy efficiency

Microsoft has operated carbon neutral since 2012. By 2030, Microsoft will be carbon negative and by 2050, Microsoft will have removed from the environment all the carbon the company has emitted either directly or by electrical consumption since it was founded in 1975. By 2025, Microsoft datacenters will be supplied by 100 percent renewable energy. To accomplish this goal, Microsoft has implemented contracting tools such as proxy generation power purchase agreements for green energy to supply 100 percent of the carbon-emitting electricity consumed by all datacenters.

Datacenter physical access security

Article • 03/02/2023

Microsoft understands the importance of protecting customer data and is committed to securing the datacenters that contain it. Microsoft datacenters are designed, built, and operated to strictly limit physical access to the areas where customer data is stored.

Physical security at datacenters is in alignment with the defense-in-depth principle. Multiple security measures are implemented to reduce the risk of unauthorized users accessing data and other datacenter resources.

- **Perimeter security**: Microsoft datacenters are nondescript buildings with perimeter fencing and 24-hour exterior lighting. Tall fences made of steel and concrete encompass every inch of the perimeter and all entry to the datacenter campus must go through a well-defined access point. Camera-monitored entrance gates and security guard patrols ensure entry and exit are restricted to designated areas. Bollards and other measures protect the datacenter exterior from potential threats, including unauthorized access.
- **Entering the datacenter**: The datacenter entrance is staffed with professional security officers who have undergone rigorous training and background checks. Security officers routinely patrol the datacenter and video feeds from cameras inside the datacenter are always monitored.
- **Inside the datacenter**: Upon entering the building, two-factor authentication with biometrics is required to continue moving through the datacenter. Once authenticated, access is granted to the authorized portion of the datacenter and only for the time approved. Within the datacenter, areas designated as highly sensitive require additional two-factor authentication.
- **Datacenter floor**: The datacenter floor can only be accessed with prior approval and after a full body metal detection screening at the time of entry. To reduce the risk of unauthorized data entering or leaving the datacenter, only approved devices can make their way onto the datacenter floor. Additionally, video cameras monitor the front and back of every server rack. When exiting the datacenter floor, all individuals are subject to an additional full body metal detection screening.
- **Leaving the datacenter**: To leave the datacenter facility, each person must go through a final security checkpoint and all visitors must surrender their temporary badges. After collection, all visitor badges have their access levels removed before they are reused for future visits.

Access provisioning

The Datacenter Management (DCM) team has implemented operational procedures to restrict physical access to only authorized employees, contractors, and visitors. Temporary or permanent access requests are tracked using a ticketing system. Badges are either issued or activated for personnel requiring access after verification of identification. Physical keys and temporary access badges are secured within the security operations center (SOC).

Microsoft datacenters are subject to a least privileged access policy, meaning datacenter access is restricted to personnel with an approved business need, with no more access than necessary. Access requests are time-limited and are only renewed if the requestor's business need remains valid.

Datacenter access records are maintained in the form of approved requests. Requests can only be approved by the DCM team, and visitor access requests to datacenters are recorded and made available for any future investigations

Datacenter security personnel

Security personnel throughout the datacenter facilities and campus are responsible for the following activities:

- Operate workstations located at the Security Operations Center within the primary administration building.
- Conduct periodic inspections through facility walkthroughs and patrols of the grounds.
- Respond to fire alarms and safety issues
- Dispatch security personnel to assist service requests and emergencies
- Provide Datacenter Management team with periodic updates about security events and entry logs
- Operate and monitor alarm, access control, and surveillance systems

Visitor access

Security verification and check-in is required for personnel requiring temporary access to the interior of the datacenter facility, including tour groups and other visitors. Datacenter visitors must sign a non-disclosure agreement, undergo datacenter management review, and obtain approval by datacenter management before their scheduled visit. Upon initial arrival, datacenter visitors are processed with temporary, least privileged access credentials. Additionally, a Microsoft full-time employee (FTE) or authorized designee approved by datacenter management is assigned to escort visitors during the visit.

All visitors that have approved access to the datacenter are designated as *Escort Only* on their badges and are required to always remain with their escorts. Escorted visitors do not have any access levels granted to them and can only travel on the access of their escorts. The escort is responsible for reviewing the actions and access of their visitor during their visit to the datacenter.

Visitor access is monitored by the assigned escort and by the Control Room Supervisor via closed circuit television (CCTV) and the alarm monitoring system. Visitors with an approved access request have their access request reviewed at the time their identification is verified against a form of government-issued identification or Microsoft-issued badge. Visitors approved for escorted access are issued a self-expiring sticky badge, and upon return the access record within the tool is terminated. If a visitor leaves with their badge, the badge automatically expires within 24 hours.

Temporary access badges are stored within the access-controlled SOC and inventoried at the beginning and end of each shift. Security officers are staffed 24x7, and physical keys are stored in an electronic key management system that is linked to the physical access system requiring a security officer's PIN and access badge to gain access.

Access review and deprovisioning

The DCM team is responsible for reviewing datacenter access regularly and for conducting a quarterly audit to verify individual access is still required.

For terminations or transfers, the person's access is immediately removed from the system and their access badge is removed. This removes any datacenter access the person may have had. DCM teams also perform quarterly access reviews to validate the appropriateness of the datacenter access list in the system.

Key management

Physical/hard keys are checked out to specific personnel by matching the person's access badge to the physical key. A person must have the appropriate access level in the tool to check out specific keys. Keys are not allowed offsite.

The hard keys and badges are kept under strict control by Microsoft and are audited daily. Microsoft also mitigates risks by implementing strict assignment of access levels, and controlled distribution and management of keys. The primary access methods at datacenters are electronic access badges and biometrics, which allows for immediate revocation of access as required. Microsoft has procedures in place to determine appropriate action commensurate with the risk for all lost keys. These actions could

require the rekeying of a single-server rack or door and up to the rekeying of the entire datacenter facility.

Access logging and monitoring

Access requests and entrance/exit events are logged and retained as part of an electronic audit trail, allowing after the fact data interrogation and reconciliation. Access control system reports and data analysis allow further anomaly detection to identify and prevent unnecessary and unauthorized access.

Datacenter surveillance systems monitor critical datacenter areas like datacenter main entry/exit, datacenter colocations entry/exit, cages, locked cabinets, aisle ways, shipping and receiving areas, critical environments, perimeter doors, and parking areas. Surveillance recordings are retained for a minimum of 90 days unless local law dictates otherwise.

A Control Room Supervisor is always in the SOC to provide monitoring of physical access in the datacenter. Video surveillance is employed to monitor physical access to the datacenter and information system. The video surveillance system is linked to the building alarm monitoring system to support physical access monitoring of alarm points. Security officers ensure that only those personnel with proper authorization are allowed access and verify that anyone bringing equipment in and out of critical infrastructure facilities follows proper procedures.

Security events that occur within the datacenter are documented by the security team in a report called a Security Event Notification (SEN). SEN reports capture the details of a security event and are required to be documented after an event occurs to capture details as accurately as possible. SEN reports also contain the investigative analysis conducted in an After Action Report (AAR), which documents the investigation of a security event, attempts to identify the root cause of the event, and records any remediation actions and lessons learned. Remediation actions and lessons learned are utilized to improve security procedures and reduce the likeliness of the event repeating. If an incident impacts Microsoft assets or services, the Security Incident Management (SIM) team has detailed procedures in place to respond.

In addition to the 24x7 onsite security, Microsoft datacenters utilize alarm monitoring systems that provide real-time alarm and video monitoring. Datacenter doors have alarms that report on each opening and when they remain open past a programmed length of time. The security system is programmed to display live video image when a door alarm is triggered. Access card and biometric readers are programmed and monitored through the alarm monitoring system. Alarms are monitored and responded

to 24x7 by the Control Room Supervisor who utilizes cameras in the area of the incident being investigated to give the responder real-time information.

Datacenter asset management

Article • 03/02/2023

Microsoft datacenters are made up of various physical and virtual components, each of which is referred to as an asset. The Microsoft Cloud Operations and Innovation (CO+I) engineering group follows secure standardized processes for the deployment, tracking, and decommissioning of both physical and virtual assets.

Supply chain integrity

Securing assets begins with securing the supply chain. Microsoft is committed to supply chain integrity and end-to-end supply chain security. Suppliers follow strict chain-of-custody procedures when transporting our cloud components to reduce the risk of alteration and tampering. All inbound and outbound inventory is carefully inspected and monitored to ensure firmware and component integrity.

Information system component deliveries to Microsoft datacenters are typically scheduled. For unscheduled deliveries, Microsoft thoroughly inspects and ensures they are approved for entry to a Microsoft computer room prior to physical entry. When an information system component enters the building, the Asset Management team compares the received item against the referenced ticket and scans the device into the asset management tool. All information system components received or shipped are tracked by the workflow ticketing tool and in the shipping logs in the asset management tool. Visitors are prohibited from using personal laptops and cell phones (except for emergency calling and taking notes), although they are allowed in the production environment (colocations). Plugging cellphones into a device or taking photos are prohibited per datacenter policy. If the equipment entering the datacenter is used for maintenance purposes, the equipment requires Datacenter Management approval in the Datacenter Access Tool system.

Asset inventory

Microsoft requires all datacenter assets to be accounted for and to have a designated owner. Owners are responsible for maintaining up-to-date asset information for their physical and virtual assets. When new physical assets are added to a datacenter, they are signed for, scanned, uniquely tagged, and checked-in to the inventory control systems. Automated monitoring tools help track both physical and virtual assets.

Asset maintenance

Microsoft has two teams that provide different types of asset maintenance: the Critical Environment (CE) and Site Services teams. The CE team provides facility management for electrical, mechanical, and physical systems that comprise the operating infrastructure of the facility. They also schedule, perform, document, and review all maintenance activities performed on CE components. Microsoft datacenters rely on a computerized maintenance management system to manage maintenance schedules and work orders. The Site Services team services all the Microsoft online service assets located at the Microsoft datacenters. Additionally, the Site Services team provides on-site technical support and break-fix service for assets belonging to properties provisioning services from the datacenter.

Asset classification

Microsoft assets - including data - are classified in accordance with Enterprise Data Taxonomy guidelines. These guidelines promote standardization across the enterprise and are reviewed and updated annually. The asset classification and asset protection standards outline the security procedures employees must follow when interacting with each asset. Customers are considered the owner of their data stored in the Microsoft cloud environment. Data assets classified as customer content or customer data are protected by applicable security procedures.

Microsoft considers all documentation to be categorized as a system asset. The Site Services team is responsible for classifying its assets and employing the associated safeguards according to the asset classification and asset protection standards, as well as any additional requirements defined by the service team.

Asset owners are required to assign their assets an asset classification, without exception. In Microsoft datacenter environments, assets refer to servers and network devices. Other digital media like USB flash drives, external hard drives, or CD/DVDs are not used for services operation. Non-digital media is not used in datacenters.

Media storage

Microsoft digital media assets are physically and securely stored within Microsoft datacenter colocation rooms. Multiple layers of physical access controls and video surveillance are in place to protect and monitor these assets. When digital media assets reach the end of their lifecycle, they are cleared, purged, or destroyed using methods

consistent with NIST SP 800-88 prior to disposal. The data-bearing device destruction article covers the process of secure disposal of assets.

Media transport

Microsoft restricts asset transport activities to authorized personnel through chain of custody protections. The use of locks, tamper-proof seals, and required validation of the asset inventory ensures that only authorized personnel are involved in the asset transport.

All media being transported from Microsoft datacenters requires accurate tracking. Microsoft has contracted with several approved vendors to provide secure shipping services. Secure transport begins with an accurate inventory and chain of custody. At the delivery location, the transport company's approved personnel must witness the removal of the tamper-proof seal and unlocking of the container. The receiving personnel inventories the shipment and sends a message confirming the receipt of the assets. Microsoft also contracts with a vendor to provide equipment destruction. Depending on the asset classification, some equipment is required to be destroyed onsite.

The transport of Microsoft digital media outside of the security-controlled space requires the supervision of two Datacenter Operations team members. The assets are continuously accompanied and supervised by authorized personnel until they are destroyed.

The Datacenter Management team controls the delivery and removal of information system components through tickets tracked in the ticketing tool. Delivery and removal of information system components is authorized by system owners.

Asset retention

Microsoft owned assets are retained as appropriate based on retention requirements set by Corporate Records Management, the asset classification, or contractual requirements. For more information about data retention, see Data retention, deletion, and destruction in Microsoft 365.

Data-bearing device destruction

Article • 03/02/2023

Data destruction overview

Microsoft has data-bearing device (DBD) guidelines, policies, security requirements, and procedures for handling and management of DBDs within Microsoft datacenters.

A DBD is any storage device capable of storing customer or proprietary Microsoft data:

- Hard disk drives (HDD)
- Solid-state drives (SSD)
- USB drives
- IO Accelerator cards
- SD/Compact Flash cards
- HSM cards
- PCIe SSD cards
- NVDIMM (Non-Volatile Dual In-line Memory Module)

Failed DBDs used within Microsoft datacenters are audited and destroyed within the datacenter campus. Any asset retired from service is evaluated for disposal in a manner commensurate with its security/privacy requirements and asset classification, and in accordance with any applicable rules, laws, and regulations.

Microsoft uses three categories of data sanitization for DBDs and assets containing data:

- **Clear**: relates to the logical techniques that help to sanitize data in all user-addressable storage locations for protection against simple non-invasive data recovery techniques. These are techniques typically applied through the standard read and write commands to the storage device, such as by rewriting with a new value or using a menu option to reset the device to the factory state (where rewriting is not supported).
- **Purge**: relates to physical or logical techniques that render target data recovery infeasible using state-of-the-art laboratory techniques.
- **Destroy**: renders target data recovery infeasible using state-of-the-art laboratory techniques and results in the subsequent inability to use the media for storage of data.

Purge and destroy sanitization is performed using tools and processes approved by the Security Group. Records are kept of the erasure and destruction of assets. Devices that

fail to complete the clear are successfully degaussed (for magnetic media only), multi-pin punched (for chipped based boards such as SSDs), or destroyed.

Clear

If a retired asset is evaluated and deemed to be non-accessible, it is cleared by an approved data eradication solution. Microsoft datacenters use the NIST SP-800-88 clear guidelines.

Purge

Depending on the on-site configuration and device availability, some devices are purged before destruction. Purge devices include NSA-approved degaussers for magnetic media and multi-pin punch devices for solid-state media. Microsoft datacenters use the NIST SP-800-88 purge guidelines.

Destroy

If a retired asset is evaluated and deemed to be accessible, it is destroyed onsite using an approved standard operating procedure that meets NIST SP-800-88 guidelines. These DBDs are physically and logically tracked to maintain chain of custody through final disposition.

Each Microsoft datacenter uses an on-site process to sanitize and dispose of failed and retired DBDs. During this process, Microsoft personnel ensure chain of custody is maintained throughout the disposal process.

Resources

- NIST Special Publication 800-88

Datacenter business continuity and disaster recovery

Article • 06/05/2023

Disasters are unpredictable, but Microsoft datacenters and operations personnel prepare for disasters to provide continuity of operations should unexpected events occur. Resilient architecture and up-to-date tested continuity plans mitigate potential damage and promote swift recovery of datacenter operations. Crisis management plans provide clarity on roles, responsibilities, and mitigation activities before, during, and after a crisis. The roles and contacts defined in these plans facilitate effective escalation up the chain of command during crisis situations.

Business resilience

Under Microsoft Cloud Operations and Innovation (CO+I) Business Continuity Program, datacenters are required to test the continued operation and response to crisis events. Each Microsoft managed datacenter has their own business continuity plan, created by using the key subject matter expertise from CO+I Resilience Center of Excellence and Datacenter Operations to ensure that site-specific context is factored into emergency preparedness. These plans describe roles, responsibilities, personnel safety procedures, notification criteria, escalations steps, and checklists for different disaster scenarios.

Microsoft's CO+I organization's Resilience function is governed by the Enterprise Business Continuity Management program and follows the Enterprise Policies and Standards. The program's performance is reviewed on a periodic basis by the Business Continuity Council, departmental leadership, and ultimately Microsoft's Senior Leadership Team.

Crisis management and pandemic response

The Crisis Management Program is an integral part of Microsoft's response to major events given its global presence. Microsoft's Datacenter Crisis Management Plan is based on industry-best practices and includes the critical components required to allow for a tactical approach in responding to major events. In addition, CO+I Resilience Center of Excellence developed and continues to maintain a Pandemic and Infectious Disease Plan that is used to respond to infectious diseases that may have an operational impact. As part of our pandemic response, the resilience support team provides critical

and timely local disease intelligence to Redmond-based Microsoft leadership to facilitate a comprehensive mitigation strategy.

Microsoft has established an organization-wide Enterprise Resilience and Crisis Management (ERCM) framework that serves as a guideline for developing Business Continuity Program across the company. The program includes Business Continuity Policy, Implementation Guidelines, Business Impact Analysis (BIA), Risk Assessment, Dependency Analysis, and procedures for monitoring and improving the program. Enterprise Resilience Office manages the governance and performance reporting across Microsoft. The CO+I Resilience program is coordinated through the CO+I Resilience Center of Excellence to ensure that the program adheres to a coherent long-term vision and mission, and is consistent with enterprise program standards, methods, policies, and metrics. CO+I Resilience Center of Excellence established a series of Standards designed to provide additional governance to the CO+I organization.

The CO+I Technology Resilience Plans (TRPs) are intended for various Engineering Groups within CO+I for the recovery from high-severity incidents or disasters to help ensure that our critical technology remains available.

The Business Resilience Plan (BRP) and TRP includes scope and applicable dependencies for the services, restoration procedures, and communications with the Incident Management team. The BRP and TRP are reviewed and approved at least annually by dedicated plan owners and made available to all applicable users. The plans are tested per the defined testing schedule as part of applicable standards.

Resiliency Program

Microsoft has defined the BRP to serve as a guide to respond, recover, and resume operations during a serious adverse event. The BRP covers the key personnel, resources, services, and actions required to continue critical business processes and operations. The development of the BRP is based on recommended guidelines of Microsoft's Enterprise Resilience Office.

In scope for this plan are Microsoft's critical business processes, defined as needed within 24 hours or less. These processes are determined during a BIA, in which Microsoft estimated potential operational and financial impacts if they could not perform a process and determined the Recovery Time Objective (RTO) and Recovery Point Objective (RPO). Following the BIA, a Non-Technical Dependency Analysis is performed to determine the specific people, applications, vital records, and user requirements necessary to perform the process.

Microsoft periodically tests the BRP to assess its effectiveness, usability, and to identify areas where risks can be eliminated or mitigated. When applicable, third parties are involved in the test if there are dependencies associated with them. The results of testing are documented, validated, approved by appropriate personnel. This information is used to create and prioritize work items.

Datacenter Resilience Program

As part of the datacenter Resilience program, the CO+I Resilience Center of Excellence team develops the methods, policies, and metrics that address the information security requirements needed for the organization's business continuity. The team develops TRPs for the continued operations of critical processes and required resources if disruptions occur.

Firmware security

Article • 11/11/2022

This article describes how Microsoft secures the cloud hardware ecosystem and supply chains.

Securing the cloud hardware ecosystem

Microsoft actively partners within the cloud hardware ecosystem to drive continuous security improvements by:

- Collaborating with Azure hardware and firmware partners (such as component manufacturers and system integrators) to meet Azure hardware and firmware security requirements.
- Enabling partners to perform continuous assessment and improvement of their products' security posture using Microsoft-defined requirements in areas such as:
 - Firmware secure boot
 - Firmware secure recovery
 - Firmware secure update
 - Firmware cryptography
 - Locked down hardware
 - Granular debug telemetry
 - System support for TPM 2.0 hardware to enable measured boot
- Engaging in and contributing to the Open Compute Project (OCP) security project through the development of specifications. Specifications promote consistency and clarity for secure design and architecture in the ecosystem.

> **Note**
>
> An example of our contribution to the OCP Security Project is the **Hardware Secure Boot** specification.

Securing hardware and firmware supply chains

Cloud hardware suppliers and vendors for Azure are also required to adhere to supply chain security processes and requirements developed by Microsoft. Hardware and

firmware development and deployment processes are required to follow the Microsoft Security Development Lifecycle (SDL) processes such as:

- Threat modeling
- Secure design reviews
- Firmware reviews and penetration testing
- Secure build and test environments
- Security vulnerability management and incident response

Next steps

To learn more about what we do to drive platform integrity and security, see:

- Platform code integrity
- Secure boot
- Measured boot and host attestation
- Project Cerberus
- Encryption at rest
- Hypervisor security

Encryption and key management overview

Article • 05/25/2023

What role does encryption play in protecting customer content?

Most Microsoft business cloud services are multi-tenant, meaning that customer content may be stored on the same physical hardware as other customers. To protect the confidentiality of customer content, Microsoft online services encrypt all data at rest and in transit with some of the strongest and most secure encryption protocols available.

Encryption isn't a substitute for strong access controls. Microsoft's access control policy of Zero Standing Access (ZSA) protects customer content from unauthorized access by Microsoft employees. Encryption complements access control by protecting the confidentiality of customer content wherever it's stored and by preventing content from being read while in transit between Microsoft online services systems or between Microsoft online services and the customer.

How do Microsoft online services encrypt data-at-rest?

All customer content in Microsoft online services is protected by one or more forms of encryption. Microsoft servers use BitLocker to encrypt the disk drives containing customer content at the volume-level. The encryption provided by BitLocker protects customer content if there are lapses in other processes or controls (for example, access control or recycling of hardware) that could lead to unauthorized physical access to disks containing customer content.

In addition to volume-level encryption, Microsoft online services use Service Encryption at the application layer to encrypt customer content. Service Encryption provides rights protection and management features on top of strong encryption protection. It also allows for separation between Windows operating systems and the customer data stored or processed by those operating systems.

How do Microsoft online services encrypt data-in-transit?

Microsoft online services use strong transport protocols, such as TLS, to prevent unauthorized parties from eavesdropping on customer data while it moves over a network. Examples of data in transit include mail messages that are in the process of being delivered, conversations taking place in an online meeting, or files being replicated between datacenters.

For Microsoft online services, data is considered 'in transit' whenever a user's device is communicating with a Microsoft server, or a Microsoft server is communicating with another server.

How do Microsoft online services manage the keys used for encryption?

Strong encryption is only as secure as the keys used to encrypt data. Microsoft uses its own security certificates to encrypt TLS connections for data-in-transit. For data-at-rest, BitLocker-protected volumes are encrypted with a full volume encryption key, which is encrypted with a volume master key, which in turn is bound to the Trusted Platform Module (TPM) in the server. BitLocker uses FIPS-compliant algorithms to ensure that encryption keys are never stored or sent over the wire in the clear.

Service Encryption provides another layer of encryption for customer data-at-rest giving customers two options for encryption key management: Microsoft-managed keys or Customer Key. When using Microsoft-managed keys, Microsoft online services automatically generate and securely store the root keys used for Service Encryption.

Customers with requirements to control their own root encryption keys can use Service Encryption with Microsoft Purview Customer Key. Using Customer Key, customers can generate their own cryptographic keys using either an on-premises Hardware Service Module (HSM) or Azure Key Vault (AKV). Customer root keys are stored in AKV, where they can be used as the root of one of the keychains that encrypts customer mailbox data or files. Customer root keys can only be accessed indirectly by Microsoft online service code for data encryption and can't be accessed directly by Microsoft employees.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to

encryption and key management.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
ISO 27001/27002 Statement of Applicability Certificate	A.10.1: Cryptographic controls A.18.1.5: Cryptographic controls	April 24, 2023
ISO 27017 Statement of Applicability Certificate	A.10.1: Cryptographic controls A.18.1.5: Cryptographic controls	April 24, 2023
ISO 27018 Statement of Applicability Certificate	A.11.6: Encryption of PII transmitted over public data transmission networks	April 24, 2023
SOC 1 SOC 2 SOC 3	DS-1: Secure storage of cryptographic certificates and keys DS-2: Customer data is encrypted in-transit DS-3: Internal communication of Azure components encrypted in-transit DS-4: Cryptographic controls and procedures	November 17, 2023

Microsoft 365

Expand table

External audits	Section	Latest report date
FedRAMP (Office 365)	SC-8: Transmission confidentiality and integrity SC-13: Use of cryptography SC-28: Protection of information at rest	July 31, 2023
ISO 27001/27002/27017	A.10.1: Cryptographic controls A.18.1.5: Cryptographic controls	March 2022

External audits	Section	Latest report date
Statement of Applicability Certification (27001/27002) Certification (27017)		
ISO 27018 Statement of Applicability Certificate	A.11.6: Encryption of PII transmitted over public data transmission networks	March 2022
SOC 2	CA-44: Data-in-transit encryption CA-54: Data-at-rest encryption CA-62: Customer Key mailbox encryption CA-63: Customer Key data deletion CA-64: Customer Key	January 3, 2023
SOC 3	CUEC-16: Customer encryption keys CUEC-17: Customer Key vault CUEC-18: Customer Key rotation	January 3, 2023

Resources

- Video: Microsoft 365 data at rest encryption

Feedback

Was this page helpful? Yes No

Azure Data Encryption at rest

Article • 11/15/2022

Microsoft Azure includes tools to safeguard data according to your company's security and compliance needs. This paper focuses on:

- How data is protected at rest across Microsoft Azure
- Discusses the various components taking part in the data protection implementation,
- Reviews pros and cons of the different key management protection approaches.

Encryption at Rest is a common security requirement. In Azure, organizations can encrypt data at rest without the risk or cost of a custom key management solution. Organizations have the option of letting Azure completely manage Encryption at Rest. Additionally, organizations have various options to closely manage encryption or encryption keys.

What is encryption at rest?

Encryption is the secure encoding of data used to protect confidentiality of data. The Encryption at Rest designs in Azure use symmetric encryption to encrypt and decrypt large amounts of data quickly according to a simple conceptual model:

- A symmetric encryption key is used to encrypt data as it is written to storage.
- The same encryption key is used to decrypt that data as it is readied for use in memory.
- Data may be partitioned, and different keys may be used for each partition.
- Keys must be stored in a secure location with identity-based access control and audit policies. Data encryption keys which are stored outside of secure locations are encrypted with a key encryption key kept in a secure location.

In practice, key management and control scenarios, as well as scale and availability assurances, require additional constructs. Microsoft Azure Encryption at Rest concepts and components are described below.

The purpose of encryption at rest

Encryption at rest provides data protection for stored data (at rest). Attacks against data at-rest include attempts to obtain physical access to the hardware on which the data is stored, and then compromise the contained data. In such an attack, a server's hard drive

may have been mishandled during maintenance allowing an attacker to remove the hard drive. Later the attacker would put the hard drive into a computer under their control to attempt to access the data.

Encryption at rest is designed to prevent the attacker from accessing the unencrypted data by ensuring the data is encrypted when on disk. If an attacker obtains a hard drive with encrypted data but not the encryption keys, the attacker must defeat the encryption to read the data. This attack is much more complex and resource consuming than accessing unencrypted data on a hard drive. For this reason, encryption at rest is highly recommended and is a high priority requirement for many organizations.

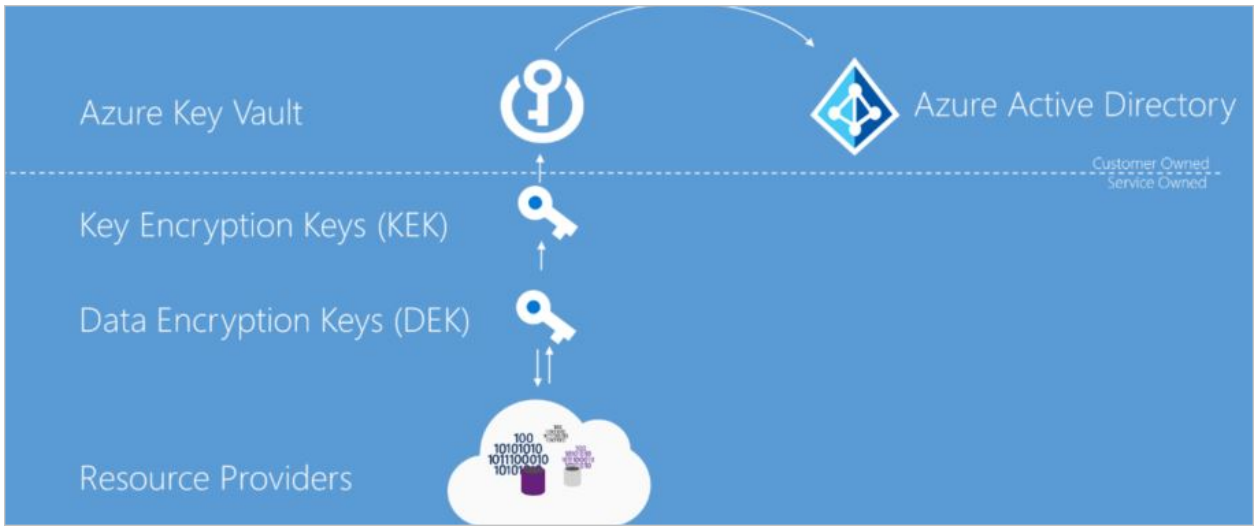
Encryption at rest may also be required by an organization's need for data governance and compliance efforts. Industry and government regulations such as HIPAA, PCI and FedRAMP, lay out specific safeguards regarding data protection and encryption requirements. Encryption at rest is a mandatory measure required for compliance with some of those regulations. For more information on Microsoft's approach to FIPS 140-2 validation, see Federal Information Processing Standard (FIPS) Publication 140-2.

In addition to satisfying compliance and regulatory requirements, encryption at rest provides defense-in-depth protection. Microsoft Azure provides a compliant platform for services, applications, and data. It also provides comprehensive facility and physical security, data access control, and auditing. However, it's important to provide additional "overlapping" security measures in case one of the other security measures fails and encryption at rest provides such a security measure.

Microsoft is committed to encryption at rest options across cloud services and giving customers control of encryption keys and logs of key use. Additionally, Microsoft is working towards encrypting all customer data at rest by default.

Azure Encryption at Rest Components

As described previously, the goal of encryption at rest is that data that is persisted on disk is encrypted with a secret encryption key. To achieve that goal secure key creation, storage, access control, and management of the encryption keys must be provided. Though details may vary, Azure services Encryption at Rest implementations can be described in terms illustrated in the following diagram.



Azure Key Vault

The storage location of the encryption keys and access control to those keys is central to an encryption at rest model. The keys need to be highly secured but manageable by specified users and available to specific services. For Azure services, Azure Key Vault is the recommended key storage solution and provides a common management experience across services. Keys are stored and managed in key vaults, and access to a key vault can be given to users or services. Azure Key Vault supports customer creation of keys or import of customer keys for use in customer-managed encryption key scenarios.

Azure Active Directory

Permissions to use the keys stored in Azure Key Vault, either to manage or to access them for Encryption at Rest encryption and decryption, can be given to Azure Active Directory accounts.

Envelope Encryption with a Key Hierarchy

More than one encryption key is used in an encryption at rest implementation. Storing an encryption key in Azure Key Vault ensures secure key access and central management of keys. However, service local access to encryption keys is more efficient for bulk encryption and decryption than interacting with Key Vault for every data operation, allowing for stronger encryption and better performance. Limiting the use of a single encryption key decreases the risk that the key will be compromised and the cost of re-encryption when a key must be replaced. Azure encryption at rest models use envelope encryption, where a key encryption key encrypts a data encryption key. This model forms a key hierarchy which is better able to address performance and security requirements:

- **Data Encryption Key (DEK)** – A symmetric AES256 key used to encrypt a partition or block of data, sometimes also referred to as simply a Data Key. A single resource may have many partitions and many Data Encryption Keys. Encrypting each block of data with a different key makes crypto analysis attacks more difficult. And keeping DEKs local to the service encrypting and decrypting data maximizes performance.
- **Key Encryption Key (KEK)** – An encryption key used to encrypt the Data Encryption Keys using envelope encryption, also referred to as wrapping. Use of a Key Encryption Key that never leaves Key Vault allows the data encryption keys themselves to be encrypted and controlled. The entity that has access to the KEK may be different than the entity that requires the DEK. An entity may broker access to the DEK to limit the access of each DEK to a specific partition. Since the KEK is required to decrypt the DEKs, customers can cryptographically erase DEKs and data by disabling of the KEK.

Resource providers and application instances store the encrypted Data Encryption Keys as metadata. Only an entity with access to the Key Encryption Key can decrypt these Data Encryption Keys. Different models of key storage are supported. For more information, see data encryption models.

Encryption at rest in Microsoft cloud services

Microsoft Cloud services are used in all three cloud models: IaaS, PaaS, SaaS. Below you have examples of how they fit on each model:

- Software services, referred to as Software as a Service or SaaS, which have applications provided by the cloud such as Microsoft 365.
- Platform services in which customers use the cloud for things like storage, analytics, and service bus functionality in their applications.
- Infrastructure services, or Infrastructure as a Service (IaaS) in which customer deploys operating systems and applications that are hosted in the cloud and possibly leveraging other cloud services.

Encryption at rest for SaaS customers

Software as a Service (SaaS) customers typically have encryption at rest enabled or available in each service. Microsoft 365 has several options for customers to verify or enable encryption at rest. For information about Microsoft 365 services, see Encryption in Microsoft 365.

Encryption at rest for PaaS customers

Platform as a Service (PaaS) customer's data typically resides in a storage service such as Blob Storage but may also be cached or stored in the application execution environment, such as a virtual machine. To see the encryption at rest options available to you, examine the Data encryption models: supporting services table for the storage and application platforms that you use.

Encryption at rest for IaaS customers

Infrastructure as a Service (IaaS) customers can have a variety of services and applications in use. IaaS services can enable encryption at rest in their Azure hosted virtual machines and VHDs using Azure Disk Encryption.

Encrypted storage

Like PaaS, IaaS solutions can leverage other Azure services that store data encrypted at rest. In these cases, you can enable the Encryption at Rest support as provided by each consumed Azure service. The Data encryption models: supporting services table enumerates the major storage, services, and application platforms and the model of Encryption at Rest supported.

Encrypted compute

All Managed Disks, Snapshots, and Images are encrypted using Storage Service Encryption using a service-managed key. A more complete Encryption at Rest solution ensures that the data is never persisted in unencrypted form. While processing the data on a virtual machine, data can be persisted to the Windows page file or Linux swap file, a crash dump, or to an application log. To ensure this data is encrypted at rest, IaaS applications can use Azure Disk Encryption on an Azure IaaS virtual machine (Windows or Linux) and virtual disk.

Custom encryption at rest

It is recommended that whenever possible, IaaS applications leverage Azure Disk Encryption and Encryption at Rest options provided by any consumed Azure services. In some cases, such as irregular encryption requirements or non-Azure based storage, a developer of an IaaS application may need to implement encryption at rest themselves. Developers of IaaS solutions can better integrate with Azure management and customer expectations by leveraging certain Azure components. Specifically, developers should

use the Azure Key Vault service to provide secure key storage as well as provide their customers with consistent key management options with that of most Azure platform services. Additionally, custom solutions should use Azure managed service identities to enable service accounts to access encryption keys. For developer information on Azure Key Vault and Managed Service Identities, see their respective SDKs.

Azure resource providers encryption model support

Microsoft Azure Services each support one or more of the encryption at rest models. For some services, however, one or more of the encryption models may not be applicable. For services that support customer-managed key scenarios, they may support only a subset of the key types that Azure Key Vault supports for key encryption keys. Additionally, services may release support for these scenarios and key types at different schedules. This section describes the encryption at rest support at the time of this writing for each of the major Azure data storage services.

Azure disk encryption

Any customer using Azure Infrastructure as a Service (IaaS) features can achieve encryption at rest for their IaaS VMs and disks through Azure Disk Encryption. For more information on Azure Disk encryption, see Azure Disk Encryption for Linux VMs or Azure Disk Encryption for Windows VMs.

Azure storage

All Azure Storage services (Blob storage, Queue storage, Table storage, and Azure Files) support server-side encryption at rest; some services additionally support customer-managed keys and client-side encryption.

- Server-side: All Azure Storage Services enable server-side encryption by default using service-managed keys, which is transparent to the application. For more information, see Azure Storage Service Encryption for Data at Rest. Azure Blob storage and Azure Files also support RSA 2048-bit customer-managed keys in Azure Key Vault. For more information, see Storage Service Encryption using customer-managed keys in Azure Key Vault.
- Client-side: Azure Blobs, Tables, and Queues support client-side encryption. When using client-side encryption, customers encrypt the data and upload the data as an encrypted blob. Key management is done by the customer. For more information, see Client-Side Encryption and Azure Key Vault for Microsoft Azure Storage.

Azure SQL Database

Azure SQL Database currently supports encryption at rest for Microsoft-managed service side and client-side encryption scenarios.

Support for server encryption is currently provided through the SQL feature called Transparent Data Encryption. Once an Azure SQL Database customer enables TDE key are automatically created and managed for them. Encryption at rest can be enabled at the database and server levels. As of June 2017, Transparent Data Encryption (TDE) is enabled by default on newly created databases. Azure SQL Database supports RSA 2048-bit customer-managed keys in Azure Key Vault. For more information, see Transparent Data Encryption with Bring Your Own Key support for Azure SQL Database and Data Warehouse.

Client-side encryption of Azure SQL Database data is supported through the Always Encrypted feature. Always Encrypted uses a key that created and stored by the client. Customers can store the master key in a Windows certificate store, Azure Key Vault, or a local Hardware Security Module. Using SQL Server Management Studio, SQL users choose what key they'd like to use to encrypt which column.

Conclusion

Protection of customer data stored within Azure Services is of paramount importance to Microsoft. All Azure hosted services are committed to providing Encryption at Rest options. Azure services support either service-managed keys, customer-managed keys, or client-side encryption. Azure services are broadly enhancing Encryption at Rest availability and new options are planned for preview and general availability in the upcoming months.

Next steps

- See data encryption models to learn more about service-managed keys and customer-managed keys.
- Learn how Azure uses double encryption to mitigate threats that come with encrypting data.
- Learn what Microsoft does to ensure platform integrity and security of hosts traversing the hardware and firmware build-out, integration, operationalization, and repair pipelines.

Azure encryption overview

Article • 02/23/2023

This article provides an overview of how encryption is used in Microsoft Azure. It covers the major areas of encryption, including encryption at rest, encryption in flight, and key management with Azure Key Vault. Each section includes links to more detailed information.

Encryption of data at rest

Data at rest includes information that resides in persistent storage on physical media, in any digital format. The media can include files on magnetic or optical media, archived data, and data backups. Microsoft Azure offers a variety of data storage solutions to meet different needs, including file, disk, blob, and table storage. Microsoft also provides encryption to protect Azure SQL Database, Azure Cosmos DB, and Azure Data Lake.

Data encryption at rest is available for services across the software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS) cloud models. This article summarizes and provides resources to help you use the Azure encryption options.

For a more detailed discussion of how data at rest is encrypted in Azure, see Azure Data Encryption-at-Rest.

Azure encryption models

Azure supports various encryption models, including server-side encryption that uses service-managed keys, customer-managed keys in Key Vault, or customer-managed keys on customer-controlled hardware. With client-side encryption, you can manage and store keys on-premises or in another secure location.

Client-side encryption

Client-side encryption is performed outside of Azure. It includes:

- Data encrypted by an application that's running in the customer's datacenter or by a service application.
- Data that is already encrypted when it is received by Azure.

With client-side encryption, cloud service providers don't have access to the encryption keys and cannot decrypt this data. You maintain complete control of the keys.

Server-side encryption

The three server-side encryption models offer different key management characteristics, which you can choose according to your requirements:

- **Service-managed keys**: Provides a combination of control and convenience with low overhead.
- **Customer-managed keys**: Gives you control over the keys, including Bring Your Own Keys (BYOK) support, or allows you to generate new ones.
- **Service-managed keys in customer-controlled hardware**: Enables you to manage keys in your proprietary repository, outside of Microsoft control. This characteristic is called Host Your Own Key (HYOK). However, configuration is complex, and most Azure services don't support this model.

Azure disk encryption

You can protect your managed disks by using Azure Disk Encryption for Linux VMs, which uses DM-Crypt, or Azure Disk Encryption for Windows VMs, which uses Windows BitLocker, to protect both operating system disks and data disks with full volume encryption.

Encryption keys and secrets are safeguarded in your Azure Key Vault subscription. By using the Azure Backup service, you can back up and restore encrypted virtual machines (VMs) that use Key Encryption Key (KEK) configuration.

Azure Storage Service Encryption

Data at rest in Azure Blob storage and Azure file shares can be encrypted in both server-side and client-side scenarios.

Azure Storage Service Encryption (SSE) can automatically encrypt data before it is stored, and it automatically decrypts the data when you retrieve it. The process is completely transparent to users. Storage Service Encryption uses 256-bit Advanced Encryption Standard (AES) encryption, which is one of the strongest block ciphers available. AES handles encryption, decryption, and key management transparently.

Client-side encryption of Azure blobs

You can perform client-side encryption of Azure blobs in various ways.

You can use the Azure Storage Client Library for .NET NuGet package to encrypt data within your client applications prior to uploading it to your Azure storage.

To learn more about and download the Azure Storage Client Library for .NET NuGet package, see Windows Azure Storage 8.3.0.

When you use client-side encryption with Key Vault, your data is encrypted using a one-time symmetric Content Encryption Key (CEK) that is generated by the Azure Storage client SDK. The CEK is encrypted using a Key Encryption Key (KEK), which can be either a symmetric key or an asymmetric key pair. You can manage it locally or store it in Key Vault. The encrypted data is then uploaded to Azure Storage.

To learn more about client-side encryption with Key Vault and get started with how-to instructions, see Tutorial: Encrypt and decrypt blobs in Azure Storage by using Key Vault.

Finally, you can also use the Azure Storage Client Library for Java to perform client-side encryption before you upload data to Azure Storage, and to decrypt the data when you download it to the client. This library also supports integration with Key Vault for storage account key management.

Encryption of data at rest with Azure SQL Database

Azure SQL Database is a general-purpose relational database service in Azure that supports structures such as relational data, JSON, spatial, and XML. SQL Database supports both server-side encryption via the Transparent Data Encryption (TDE) feature and client-side encryption via the Always Encrypted feature.

Transparent Data Encryption

TDE is used to encrypt SQL Server, Azure SQL Database, and Azure Synapse Analytics data files in real time, using a Database Encryption Key (DEK), which is stored in the database boot record for availability during recovery.

TDE protects data and log files, using AES and Triple Data Encryption Standard (3DES) encryption algorithms. Encryption of the database file is performed at the page level. The pages in an encrypted database are encrypted before they are written to disk and are decrypted when they're read into memory. TDE is now enabled by default on newly created Azure SQL databases.

Always Encrypted feature

With the Always Encrypted feature in Azure SQL you can encrypt data within client applications prior to storing it in Azure SQL Database. You can also enable delegation of on-premises database administration to third parties and maintain separation between those who own and can view the data and those who manage it but should not have access to it.

Cell-level or column-level encryption

With Azure SQL Database, you can apply symmetric encryption to a column of data by using Transact-SQL. This approach is called cell-level encryption or column-level encryption (CLE), because you can use it to encrypt specific columns or even specific cells of data with different encryption keys. Doing so gives you more granular encryption capability than TDE, which encrypts data in pages.

CLE has built-in functions that you can use to encrypt data by using either symmetric or asymmetric keys, the public key of a certificate, or a passphrase using 3DES.

Azure Cosmos DB database encryption

Azure Cosmos DB is Microsoft's globally distributed, multi-model database. User data that's stored in Azure Cosmos DB in non-volatile storage (solid-state drives) is encrypted by default. There are no controls to turn it on or off. Encryption at rest is implemented by using a number of security technologies, including secure key storage systems, encrypted networks, and cryptographic APIs. Encryption keys are managed by Microsoft and are rotated per Microsoft internal guidelines. Optionally, you can choose to add a second layer of encryption with keys you manage using the customer-managed keys or CMK feature.

At-rest encryption in Data Lake

Azure Data Lake is an enterprise-wide repository of every type of data collected in a single place prior to any formal definition of requirements or schema. Data Lake Store supports "on by default," transparent encryption of data at rest, which is set up during the creation of your account. By default, Azure Data Lake Store manages the keys for you, but you have the option to manage them yourself.

Three types of keys are used in encrypting and decrypting data: the Master Encryption Key (MEK), Data Encryption Key (DEK), and Block Encryption Key (BEK). The MEK is used to encrypt the DEK, which is stored on persistent media, and the BEK is derived from the DEK and the data block. If you are managing your own keys, you can rotate the MEK.

Encryption of data in transit

Azure offers many mechanisms for keeping data private as it moves from one location to another.

Data-link Layer encryption in Azure

Whenever Azure Customer traffic moves between datacenters-- outside physical boundaries not controlled by Microsoft (or on behalf of Microsoft)-- a data-link layer encryption method using the IEEE 802.1AE MAC Security Standards (also known as MACsec) is applied from point-to-point across the underlying network hardware. The packets are encrypted on the devices before being sent, preventing physical "man-in-the-middle" or snooping/wiretapping attacks. Because this technology is integrated on the network hardware itself, it provides line rate encryption on the network hardware with no measurable link latency increase. This MACsec encryption is on by default for all Azure traffic traveling within a region or between regions, and no action is required on customers' part to enable.

TLS encryption in Azure

Microsoft gives customers the ability to use Transport Layer Security (TLS) protocol to protect data when it's traveling between the cloud services and customers. Microsoft datacenters negotiate a TLS connection with client systems that connect to Azure services. TLS provides strong authentication, message privacy, and integrity (enabling detection of message tampering, interception, and forgery), interoperability, algorithm flexibility, and ease of deployment and use.

Perfect Forward Secrecy (PFS) protects connections between customers' client systems and Microsoft cloud services by unique keys. Connections also use RSA-based 2,048-bit encryption key lengths. This combination makes it difficult for someone to intercept and access data that is in transit.

Azure Storage transactions

When you interact with Azure Storage through the Azure portal, all transactions take place over HTTPS. You can also use the Storage REST API over HTTPS to interact with Azure Storage. You can enforce the use of HTTPS when you call the REST APIs to access objects in storage accounts by enabling the secure transfer that's required for the storage account.

Shared Access Signatures (SAS), which can be used to delegate access to Azure Storage objects, include an option to specify that only the HTTPS protocol can be used when you use Shared Access Signatures. This approach ensures that anybody who sends links with SAS tokens uses the proper protocol.

SMB 3.0, which used to access Azure Files shares, supports encryption, and it's available in Windows Server 2012 R2, Windows 8, Windows 8.1, and Windows 10. It allows cross-region access and even access on the desktop.

Client-side encryption encrypts the data before it's sent to your Azure Storage instance, so that it's encrypted as it travels across the network.

SMB encryption over Azure virtual networks

By using SMB 3.0 in VMs that are running Windows Server 2012 or later, you can make data transfers secure by encrypting data in transit over Azure Virtual Networks. By encrypting data, you help protect against tampering and eavesdropping attacks. Administrators can enable SMB encryption for the entire server, or just specific shares.

By default, after SMB encryption is turned on for a share or server, only SMB 3.0 clients are allowed to access the encrypted shares.

In-transit encryption in VMs

Data in transit to, from, and between VMs that are running Windows can be encrypted in a number of ways, depending on the nature of the connection.

RDP sessions

You can connect and sign in to a VM by using the Remote Desktop Protocol (RDP) from a Windows client computer, or from a Mac with an RDP client installed. Data in transit over the network in RDP sessions can be protected by TLS.

You can also use Remote Desktop to connect to a Linux VM in Azure.

Secure access to Linux VMs with SSH

For remote management, you can use Secure Shell (SSH) to connect to Linux VMs running in Azure. SSH is an encrypted connection protocol that allows secure sign-ins over unsecured connections. It is the default connection protocol for Linux VMs hosted

in Azure. By using SSH keys for authentication, you eliminate the need for passwords to sign in. SSH uses a public/private key pair (asymmetric encryption) for authentication.

Azure VPN encryption

You can connect to Azure through a virtual private network that creates a secure tunnel to protect the privacy of the data being sent across the network.

Azure VPN gateways

You can use an Azure VPN gateway to send encrypted traffic between your virtual network and your on-premises location across a public connection, or to send traffic between virtual networks.

Site-to-site VPNs use IPsec for transport encryption. Azure VPN gateways use a set of default proposals. You can configure Azure VPN gateways to use a custom IPsec/IKE policy with specific cryptographic algorithms and key strengths, rather than the Azure default policy sets.

Point-to-site VPNs

Point-to-site VPNs allow individual client computers access to an Azure virtual network. The Secure Socket Tunneling Protocol (SSTP) is used to create the VPN tunnel. It can traverse firewalls (the tunnel appears as an HTTPS connection). You can use your own internal public key infrastructure (PKI) root certificate authority (CA) for point-to-site connectivity.

You can configure a point-to-site VPN connection to a virtual network by using the Azure portal with certificate authentication or PowerShell.

To learn more about point-to-site VPN connections to Azure virtual networks, see:

Configure a point-to-site connection to a virtual network by using certification authentication: Azure portal

Configure a point-to-site connection to a virtual network by using certificate authentication: PowerShell

Site-to-site VPNs

You can use a site-to-site VPN gateway connection to connect your on-premises network to an Azure virtual network over an IPsec/IKE (IKEv1 or IKEv2) VPN tunnel. This

type of connection requires an on-premises VPN device that has an external-facing public IP address assigned to it.

You can configure a site-to-site VPN connection to a virtual network by using the Azure portal, PowerShell, or Azure CLI.

For more information, see:

Create a site-to-site connection in the Azure portal

Create a site-to-site connection in PowerShell

Create a virtual network with a site-to-site VPN connection by using CLI

In-transit encryption in Data Lake

Data in transit (also known as data in motion) is also always encrypted in Data Lake Store. In addition to encrypting data prior to storing it in persistent media, the data is also always secured in transit by using HTTPS. HTTPS is the only protocol that is supported for the Data Lake Store REST interfaces.

To learn more about encryption of data in transit in Data Lake, see Encryption of data in Data Lake Store.

Key management with Key Vault

Without proper protection and management of the keys, encryption is rendered useless. Key Vault is the Microsoft-recommended solution for managing and controlling access to encryption keys used by cloud services. Permissions to access keys can be assigned to services or to users through Azure Active Directory accounts.

Key Vault relieves organizations of the need to configure, patch, and maintain hardware security modules (HSMs) and key management software. When you use Key Vault, you maintain control. Microsoft never sees your keys, and applications don't have direct access to them. You can also import or generate keys in HSMs.

Next steps

- Azure security overview
- Azure network security overview
- Azure database security overview
- Azure virtual machines security overview

- Data encryption at rest
- Data security and encryption best practices

Data encryption models

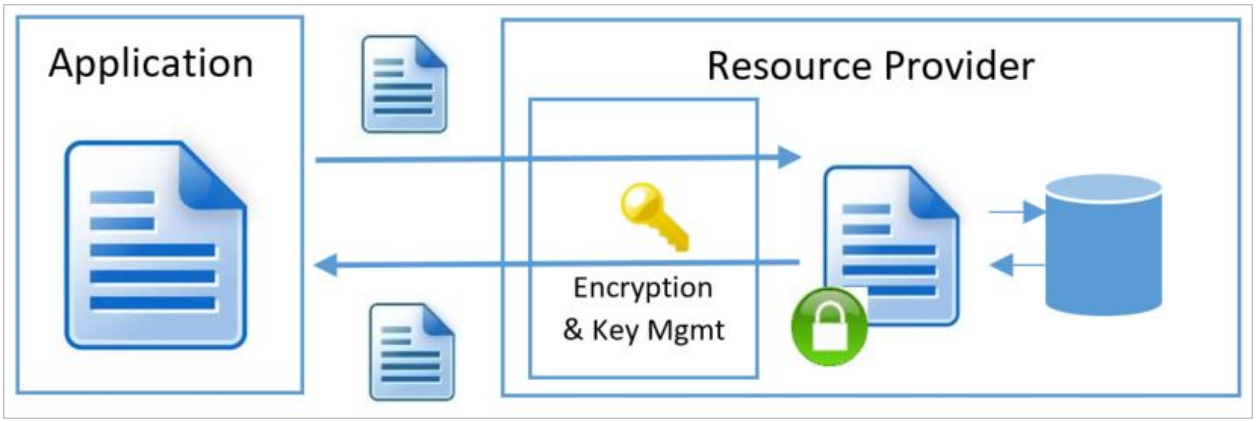
Article • 05/05/2023

An understanding of the various encryption models and their pros and cons is essential for understanding how the various resource providers in Azure implement encryption at Rest. These definitions are shared across all resource providers in Azure to ensure common language and taxonomy.

There are three scenarios for server-side encryption:

- Server-side encryption using Service-Managed keys
 - Azure Resource Providers perform the encryption and decryption operations
 - Microsoft manages the keys
 - Full cloud functionality
- Server-side encryption using customer-managed keys in Azure Key Vault
 - Azure Resource Providers perform the encryption and decryption operations
 - Customer controls keys via Azure Key Vault
 - Full cloud functionality
- Server-side encryption using customer-managed keys on customer-controlled hardware
 - Azure Resource Providers perform the encryption and decryption operations
 - Customer controls keys on customer-controlled hardware
 - Full cloud functionality

Server-side Encryption models refer to encryption that is performed by the Azure service. In that model, the Resource Provider performs the encrypt and decrypt operations. For example, Azure Storage may receive data in plain text operations and will perform the encryption and decryption internally. The Resource Provider might use encryption keys that are managed by Microsoft or by the customer depending on the provided configuration.



Each of the server-side encryption at rest models implies distinctive characteristics of key management. This includes where and how encryption keys are created, and stored as well as the access models and the key rotation procedures.

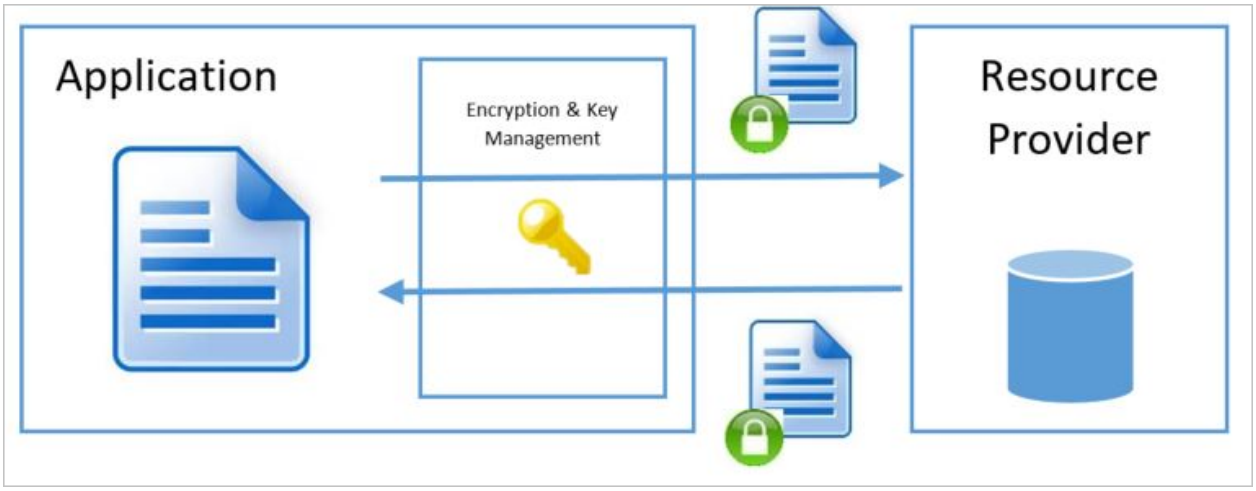
For client-side encryption, consider the following:

- Azure services cannot see decrypted data
- Customers manage and store keys on-premises (or in other secure stores). Keys are not available to Azure services
- Reduced cloud functionality

The supported encryption models in Azure split into two main groups: "Client Encryption" and "Server-side Encryption" as mentioned previously. Independent of the encryption at rest model used, Azure services always recommend the use of a secure transport such as TLS or HTTPS. Therefore, encryption in transport should be addressed by the transport protocol and should not be a major factor in determining which encryption at rest model to use.

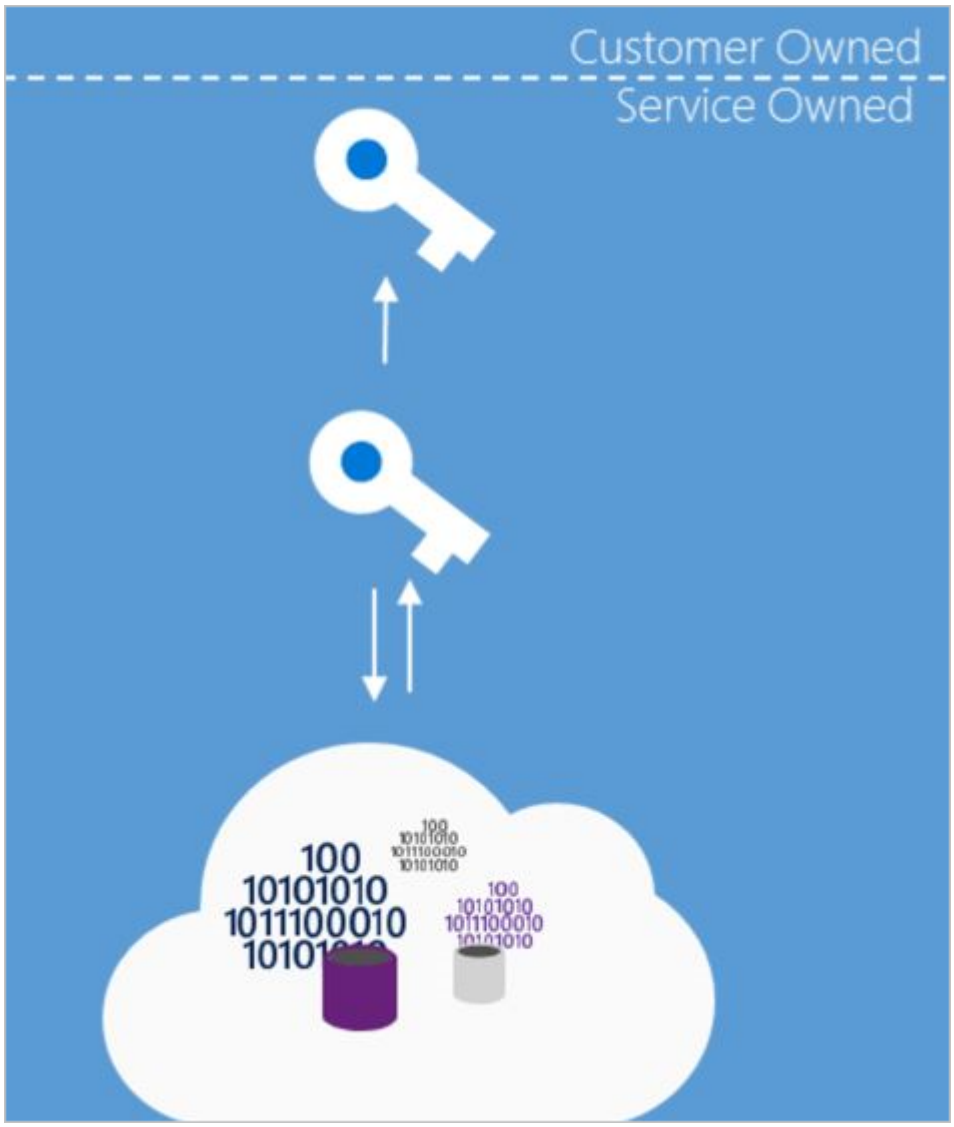
Client encryption model

Client Encryption model refers to encryption that is performed outside of the Resource Provider or Azure by the service or calling application. The encryption can be performed by the service application in Azure, or by an application running in the customer data center. In either case, when leveraging this encryption model, the Azure Resource Provider receives an encrypted blob of data without the ability to decrypt the data in any way or have access to the encryption keys. In this model, the key management is done by the calling service/application and is opaque to the Azure service.



Server-side encryption using service-managed keys

For many customers, the essential requirement is to ensure that the data is encrypted whenever it is at rest. Server-side encryption using service-managed Keys enables this model by allowing customers to mark the specific resource (Storage Account, SQL DB, etc.) for encryption and leaving all key management aspects such as key issuance, rotation, and backup to Microsoft. Most Azure services that support encryption at rest typically support this model of offloading the management of the encryption keys to Azure. The Azure resource provider creates the keys, places them in secure storage, and retrieves them when needed. This means that the service has full access to the keys and the service has full control over the credential lifecycle management.



Server-side encryption using service-managed keys therefore quickly addresses the need to have encryption at rest with low overhead to the customer. When available a customer typically opens the Azure portal for the target subscription and resource provider and checks a box indicating, they would like the data to be encrypted. In some Resource Managers server-side encryption with service-managed keys is on by default.

Server-side encryption with Microsoft-managed keys does imply the service has full access to store and manage the keys. While some customers may want to manage the keys because they feel they gain greater security, the cost and risk associated with a

custom key storage solution should be considered when evaluating this model. In many cases, an organization may determine that resource constraints or risks of an on-premises solution may be greater than the risk of cloud management of the encryption at rest keys. However, this model might not be sufficient for organizations that have requirements to control the creation or lifecycle of the encryption keys or to have different personnel manage a service's encryption keys than those managing the service (that is, segregation of key management from the overall management model for the service).

Key access

When Server-side encryption with service-managed keys is used, the key creation, storage, and service access are all managed by the service. Typically, the foundational Azure resource providers will store the Data Encryption Keys in a store that is close to the data and quickly available and accessible while the Key Encryption Keys are stored in a secure internal store.

Advantages

- Simple setup
- Microsoft manages key rotation, backup, and redundancy
- Customer does not have the cost associated with implementation or the risk of a custom key management scheme.

Disadvantages

- No customer control over the encryption keys (key specification, lifecycle, revocation, etc.)
- No ability to segregate key management from overall management model for the service

Server-side encryption using customer-managed keys in Azure Key Vault

For scenarios where the requirement is to encrypt the data at rest and control the encryption keys customers can use server-side encryption using customer-managed Keys in Key Vault. Some services may store only the root Key Encryption Key in Azure Key Vault and store the encrypted Data Encryption Key in an internal location closer to the data. In that scenario customers can bring their own keys to Key Vault (BYOK – Bring Your Own Key), or generate new ones, and use them to encrypt the desired resources.

While the Resource Provider performs the encryption and decryption operations, it uses the configured key encryption key as the root key for all encryption operations.

Loss of key encryption keys means loss of data. For this reason, keys should not be deleted. Keys should be backed up whenever created or rotated. Soft-Delete and purge protection must be enabled on any vault storing key encryption keys to protect against accidental or malicious cryptographic erasure. Instead of deleting a key, it is recommended to set enabled to false on the key encryption key. Use access controls to revoke access to individual users or services in Azure Key Vault or Managed HSM.

Key Access

The server-side encryption model with customer-managed keys in Azure Key Vault involves the service accessing the keys to encrypt and decrypt as needed. Encryption at rest keys are made accessible to a service through an access control policy. This policy grants the service identity access to receive the key. An Azure service running on behalf of an associated subscription can be configured with an identity in that subscription. The service can perform Azure Active Directory authentication and receive an authentication token identifying itself as that service acting on behalf of the subscription. That token can then be presented to Key Vault to obtain a key it has been given access to.

For operations using encryption keys, a service identity can be granted access to any of the following operations: decrypt, encrypt, unwrapKey, wrapKey, verify, sign, get, list, update, create, import, delete, backup, and restore.

To obtain a key for use in encrypting or decrypting data at rest the service identity that the Resource Manager service instance will run as must have UnwrapKey (to get the key for decryption) and WrapKey (to insert a key into key vault when creating a new key).

> ⓘ **Note**
>
> For more detail on Key Vault authorization see the secure your key vault page in the **Azure Key Vault documentation**.

Advantages

- Full control over the keys used – encryption keys are managed in the customer's Key Vault under the customer's control.
- Ability to encrypt multiple services to one master
- Can segregate key management from overall management model for the service
- Can define service and key location across regions

Disadvantages

- Customer has full responsibility for key access management
- Customer has full responsibility for key lifecycle management
- Additional Setup & configuration overhead

Server-side encryption using customer-managed keys in customer-controlled hardware

Some Azure services enable the Host Your Own Key (HYOK) key management model. This management mode is useful in scenarios where there is a need to encrypt the data at rest and manage the keys in a proprietary repository outside of Microsoft's control. In this model, the service must use the key from an external site to decrypt the Data Encryption Key (DEK). Performance and availability guarantees are impacted, and configuration is more complex. Additionally, since the service does have access to the DEK during the encryption and decryption operations the overall security guarantees of this model are similar to when the keys are customer-managed in Azure Key Vault. As a result, this model is not appropriate for most organizations unless they have specific key management requirements. Due to these limitations, most Azure services do not support server-side encryption using customer-managed keys in customer-controlled hardware. One of two keys in Double Key Encryption follows this model.

Key Access

When server-side encryption using customer-managed keys in customer-controlled hardware is used, the key encryption keys are maintained on a system configured by the customer. Azure services that support this model provide a means of establishing a secure connection to a customer supplied key store.

Advantages

- Full control over the root key used – encryption keys are managed by a customer provided store
- Ability to encrypt multiple services to one master
- Can segregate key management from overall management model for the service
- Can define service and key location across regions

Disadvantages

- Full responsibility for key storage, security, performance, and availability
- Full responsibility for key access management

- Full responsibility for key lifecycle management
- Significant setup, configuration, and ongoing maintenance costs
- Increased dependency on network availability between the customer datacenter and Azure datacenters.

Supporting services

The Azure services that support each encryption model:

Product, Feature, or Service	Server-Side Using Service-Managed Key	Server-Side Using Customer-Managed Key	Client-Side Using Client-Managed Key
AI and Machine Learning			
Azure Cognitive Search	Yes	Yes	-
Azure Cognitive Services	Yes	Yes, including Managed HSM	-
Azure Machine Learning	Yes	Yes	-
Content Moderator	Yes	Yes, including Managed HSM	-
Face	Yes	Yes, including Managed HSM	-
Language Understanding	Yes	Yes, including Managed HSM	-
Azure OpenAI	Yes	Yes, including Managed HSM	-
Personalizer	Yes	Yes, including Managed HSM	-
QnA Maker	Yes	Yes, including Managed HSM	-
Speech Services	Yes	Yes, including Managed HSM	-
Translator Text	Yes	Yes, including Managed HSM	-
Power BI	Yes	Yes, RSA 4096-bit	-
Analytics			

Product, Feature, or Service	Server-Side Using Service-Managed Key	Server-Side Using Customer-Managed Key	Client-Side Using Client-Managed Key
Azure Stream Analytics	Yes	Yes**, including Managed HSM	-
Event Hubs	Yes	Yes	-
Functions	Yes	Yes	-
Azure Analysis Services	Yes	-	-
Azure Data Catalog	Yes	-	-
Azure HDInsight	Yes	All	-
Azure Monitor Application Insights	Yes	Yes	-
Azure Monitor Log Analytics	Yes	Yes, including Managed HSM	-
Azure Data Explorer	Yes	Yes	-
Azure Data Factory	Yes	Yes, including Managed HSM	-
Azure Data Lake Store	Yes	Yes, RSA 2048-bit	-
Containers			
Azure Kubernetes Service	Yes	Yes, including Managed HSM	-
Container Instances	Yes	Yes	-
Container Registry	Yes	Yes	-
Compute			
Virtual Machines	Yes	Yes, including Managed HSM	-
Virtual Machine Scale Set	Yes	Yes, including Managed HSM	-
SAP HANA	Yes	Yes	-
App Service	Yes	Yes**, including Managed HSM	-
Automation	Yes	Yes	-

Product, Feature, or Service	Server-Side Using Service-Managed Key	Server-Side Using Customer-Managed Key	Client-Side Using Client-Managed Key
Azure Functions	Yes	Yes**, including Managed HSM	-
Azure portal	Yes	Yes**, including Managed HSM	-
Azure VMware Solution	Yes	Yes, including Managed HSM	-
Logic Apps	Yes	Yes	-
Azure-managed applications	Yes	Yes**, including Managed HSM	-
Service Bus	Yes	Yes	-
Site Recovery	Yes	Yes	-
Databases			
SQL Server on Virtual Machines	Yes	Yes	Yes
Azure SQL Database	Yes	Yes, RSA 3072-bit, including Managed HSM	Yes
Azure SQL Managed Instance	Yes	Yes, RSA 3072-bit, including Managed HSM	Yes
Azure SQL Database for MariaDB	Yes	-	-
Azure SQL Database for MySQL	Yes	Yes	-
Azure SQL Database for PostgreSQL	Yes	Yes, including Managed HSM	-
Azure Synapse Analytics (dedicated SQL pool (formerly SQL DW) only)	Yes	Yes, RSA 3072-bit, including Managed HSM	-
SQL Server Stretch Database	Yes	Yes, RSA 3072-bit	Yes
Table Storage	Yes	Yes	Yes

Product, Feature, or Service	Server-Side Using Service-Managed Key	Server-Side Using Customer-Managed Key	Client-Side Using Client-Managed Key
Azure Cosmos DB	Yes (learn more)	Yes, including Managed HSM (learn more and learn more)	-
Azure Databricks	Yes	Yes	-
Azure Database Migration Service	Yes	N/A*	-
Identity			
Azure Active Directory	Yes	-	-
Azure Active Directory Domain Services	Yes	Yes	-
Integration			
Service Bus	Yes	Yes	Yes
Event Grid	Yes	-	-
API Management	Yes	-	-
IoT Services			
IoT Hub	Yes	Yes	Yes
IoT Hub Device Provisioning	Yes	Yes	-
Management and Governance			
Azure Managed Grafana	Yes	-	N/A
Azure Site Recovery	Yes	-	-
Azure Migrate	Yes	Yes	-
Media			
Media Services	Yes	Yes	Yes
Security			
Microsoft Defender for IoT	Yes	Yes	-
Microsoft Sentinel	Yes	Yes, including Managed HSM	-

Product, Feature, or Service	Server-Side Using Service-Managed Key	Server-Side Using Customer-Managed Key	Client-Side Using Client-Managed Key
Storage			
Blob Storage	Yes	Yes, including Managed HSM	Yes
Premium Blob Storage	Yes	Yes, including Managed HSM	Yes
Disk Storage	Yes	Yes, including Managed HSM	-
Ultra Disk Storage	Yes	Yes, including Managed HSM	-
Managed Disk Storage	Yes	Yes, including Managed HSM	-
File Storage	Yes	Yes, including Managed HSM	-
File Premium Storage	Yes	Yes, including Managed HSM	-
File Sync	Yes	Yes, including Managed HSM	-
Queue Storage	Yes	Yes, including Managed HSM	Yes
Data Lake Storage Gen2	Yes	Yes, including Managed HSM	Yes
Avere vFXT	Yes	-	-
Azure Cache for Redis	Yes	N/A*	-
Azure NetApp Files	Yes	Yes	Yes
Archive Storage	Yes	Yes	-
StorSimple	Yes	Yes	Yes
Azure Backup	Yes	Yes	Yes
Data Box	Yes	-	Yes
Data Box Edge	Yes	Yes	-

Product, Feature, or Service	Server-Side Using Service-Managed Key	Server-Side Using Customer-Managed Key	Client-Side Using Client-Managed Key
Other			
Azure Data Manager for Energy Preview	Yes	-	Yes

* This service doesn't persist data. Transient caches, if any, are encrypted with a Microsoft key.

** This service supports storing data in your own Key Vault, Storage Account, or other data persisting service that already supports Server-Side Encryption with Customer-Managed Key.

Next steps

- Learn how encryption is used in Azure.
- Learn how Azure uses double encryption to mitigate threats that come with encrypting data.

Double encryption

Article • 07/03/2022

Double encryption is where two or more independent layers of encryption are enabled to protect against compromises of any one layer of encryption. Using two layers of encryption mitigates threats that come with encrypting data. For example:

- Configuration errors in the data encryption
- Implementation errors in the encryption algorithm
- Compromise of a single encryption key

Azure provides double encryption for data at rest and data in transit.

Data at rest

Microsoft's approach to enabling two layers of encryption for data at rest is:

- **Encryption at rest using customer-managed keys**. You provide your own key for data encryption at rest. You can bring your own keys to your Key Vault (BYOK – Bring Your Own Key), or generate new keys in Azure Key Vault to encrypt the desired resources.
- **Infrastructure encryption using platform-managed keys**. By default, data is automatically encrypted at rest using platform-managed encryption keys.

Data in transit

Microsoft's approach to enabling two layers of encryption for data in transit is:

- **Transit encryption using Transport Layer Security (TLS) 1.2 to protect data when it's traveling between the cloud services and you**. All traffic leaving a datacenter is encrypted in transit, even if the traffic destination is another domain controller in the same region. TLS 1.2 is the default security protocol used. TLS provides strong authentication, message privacy, and integrity (enabling detection of message tampering, interception, and forgery), interoperability, algorithm flexibility, and ease of deployment and use.
- **Additional layer of encryption provided at the infrastructure layer**. Whenever Azure customer traffic moves between datacenters-- outside physical boundaries not controlled by Microsoft or on behalf of Microsoft-- a data-link layer encryption method using the IEEE 802.1AE MAC Security Standards (also known as MACsec) is applied from point-to-point across the underlying network hardware. The

packets are encrypted and decrypted on the devices before being sent, preventing physical “man-in-the-middle” or snooping/wiretapping attacks. Because this technology is integrated on the network hardware itself, it provides line rate encryption on the network hardware with no measurable link latency increase. This MACsec encryption is on by default for all Azure traffic traveling within a region or between regions, and no action is required on customers’ part to enable.

Next steps

Learn how encryption is used in Azure.

Key management in Azure

Article • 06/29/2023

> **Note**
>
> **Zero Trust** is a security strategy comprising three principles: "Verify explicitly", "Use least privilege access", and "Assume breach". Data protection, including key management, supports the "use least privilege access" principle. For more information, see **What is Zero Trust?**

In Azure, encryption keys can be either platform managed or customer managed.

Platform-managed keys (PMKs) are encryption keys generated, stored, and managed entirely by Azure. Customers do not interact with PMKs. The keys used for Azure Data Encryption-at-Rest, for instance, are PMKs by default.

Customer-managed keys (CMK), on the other hand, are keys read, created, deleted, updated, and/or administered by one or more customers. Keys stored in a customer-owned key vault or hardware security module (HSM) are CMKs. Bring Your Own Key (BYOK) is a CMK scenario in which a customer imports (brings) keys from an outside storage location into an Azure key management service (see the Azure Key Vault: Bring your own key specification).

A specific type of customer-managed key is the "key encryption key" (KEK). A KEK is a primary key that controls access to one or more encryption keys that are themselves encrypted.

Customer-managed keys can be stored on-premises or, more commonly, in a cloud key management service.

Azure key management services

Azure offers several options for storing and managing your keys in the cloud, including Azure Key Vault, Azure Managed HSM, Azure Dedicated HSM, and Azure Payment HSM. These options differ in terms of their FIPS compliance level, management overhead, and intended applications.

For an overview of each key management service and a comprehensive guide to choosing the right key management solution for you, see How to Choose the Right Key Management Solution.

Pricing

The Azure Key Vault Standard and Premium tiers are billed on a transactional basis, with an additional monthly per-key charge for premium hardware-backed keys. Managed HSM, Dedicated HSM, and Payments HSM don't charge on a transactional basis; instead they are always-in-use devices that are billed at a fixed hourly rate. For detailed pricing information, see Key Vault pricing, Dedicated HSM pricing, and Payment HSM pricing.

Service Limits

Managed HSM, Dedicated HSM, and Payments HSM offer dedicated capacity. Key Vault Standard and Premium are multi-tenant offerings and have throttling limits. For service limits, see Key Vault service limits.

Encryption-At-Rest

Azure Key Vault and Azure Key Vault Managed HSM have integrations with Azure Services and Microsoft 365 for Customer Managed Keys, meaning customers may use their own keys in Azure Key Vault and Azure Key Managed HSM for encryption-at-rest of data stored in these services. Dedicated HSM and Payments HSM are Infrastructure-as-Service offerings and do not offer integrations with Azure Services. For an overview of encryption-at-rest with Azure Key Vault and Managed HSM, see Azure Data Encryption-at-Rest.

APIs

Dedicated HSM and Payments HSM support the PKCS#11, JCE/JCA, and KSP/CNG APIs, but Azure Key Vault and Managed HSM do not. Azure Key Vault and Managed HSM use the Azure Key Vault REST API and offer SDK support. For more information on the Azure Key Vault API, see Azure Key Vault REST API Reference.

What's next

- How to Choose the Right Key Management Solution
- Azure Key Vault
- Azure Managed HSM
- Azure Dedicated HSM
- Azure Payment HSM
- What is Zero Trust?

Control your data in the cloud by using Managed HSM

Article • 08/03/2023

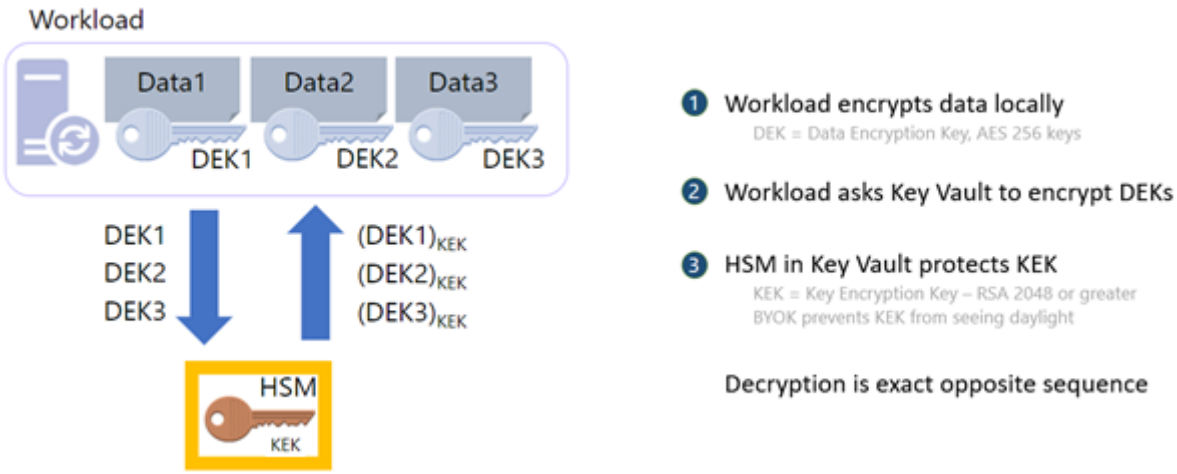
Microsoft values, protects, and defends privacy. We believe in transparency so that people and organizations can control their data and have meaningful choices in how it's used. We empower and defend the privacy choices of every person who uses our products and services.

In this article, we take a deep dive on Azure Key Vault Managed HSM security controls for encryption and how it provides additional safeguards and technical measures to help our customers meet their compliance requirements.

Encryption is one of the key technical measures you can take to get sole control of your data. Azure fortifies your data through state-of-the-art encryption technologies, both for data at rest and for data in transit. Our encryption products erect barriers against unauthorized access to the data, including two or more independent encryption layers to protect against compromises of any single layer. In addition, Azure has clearly defined, well-established responses, policies and processes, strong contractual commitments, and strict physical, operational, and infrastructure security controls to provide our customers the ultimate control of their data in the cloud. The fundamental premise of Azure's key management strategy is to give our customers more control over their data. We use a zero trust posture with advanced enclave technologies, hardware security modules (HSMs), and identity isolation that reduces Microsoft access to customer keys and data.

Encryption at rest provides data protection for stored data at rest and as required by an organization's need for data governance and compliance efforts. Microsoft's compliance portfolio is the broadest in all public clouds worldwide, with industry standards and government regulations such as HIPAA, General Data Protection Regulation, and Federal Information Processing Standards (FIPS) 140-2 and 3. These standards and regulations lay out specific safeguards for data protection and encryption requirements. In most cases, a mandatory measure is required for compliance.

How does encryption at rest work?



Azure Key Vault services provide encryption and key management solutions that safeguard cryptographic keys, certificates, and other secrets that cloud applications and services use to protect and control data that's encrypted at rest.

Secure key management is essential to protect and control data in the cloud. Azure offers various solutions that you can use to manage and control access to encryption keys so that you have choice and flexibility to meet stringent data protection and compliance needs.

- **Azure platform encryption** is a *platform-managed* encryption solution that encrypts by using host-level encryption. Platform-managed keys are encryption keys that are generated, stored, and managed entirely by Azure.
- **Customer-managed keys** are keys that are created, read, deleted, updated, and administered entirely by the customer. Customer-managed keys can be stored in a cloud key management service like Azure Key Vault.
- **Azure Key Vault Standard** encrypts by using a software key and is FIPS 140-2 Level 1 compliant.
- **Azure Key Vault Premium** encrypts by using FIPS 140-2 Level 2 HSM-protected keys.
- **Azure Key Vault Managed HSM** encrypts by using single-tenant FIPS 140-2 Level 3 HSM protected keys and is fully managed by Microsoft. It provides customers with sole control of the cryptographic keys.

For added assurance, in Azure Key Vault Premium and Azure Key Vault Managed HSM, you can bring your own key (BYOK) and import HSM-protected keys from an on-premises HSM.

Portfolio of Azure key management products

	Azure Key Vault Standard	Azure Key Vault Premium	Azure Key Vault Managed HSM
Tenancy	Multitenant	Multitenant	Single-tenant
Compliance	FIPS 140-2 Level 1	FIPS 140-2 Level 2	FIPS 140-2 Level 3
High availability	Automatic	Automatic	Automatic
Use cases	Encryption at rest	Encryption at rest	Encryption at rest
Key controls	Customer	Customer	Customer
Root of trust control	Microsoft	Microsoft	Customer

Azure Key Vault is a cloud service that you can use to securely store and access secrets. A secret is anything that you want to securely control access to and can include API keys, passwords, certificates, and cryptographic keys.

Key Vault supports two types of containers:

- Vaults
 - Standard tier: Vaults support storing secrets, certificates, and software-backed keys.
 - Premium tier: Vaults support storing secrets, certificates, software-backed keys, and HSM-backed keys.
- Managed hardware security module (HSM)
 - Managed HSM supports only HSM-backed keys.

For more information, see Azure Key Vault Concepts and Azure Key Vault REST API overview.

What is Azure Key Vault Managed HSM?

Azure Key Vault Managed HSM is a fully managed, highly available, single-tenant, standards-compliant cloud service that has a customer-controlled security domain that enables you to store cryptographic keys for your cloud applications by using FIPS 140-2 Level 3 validated HSMs.

How does Azure Key Vault Managed HSM protect your keys?

Azure Key Vault Managed HSM uses a defense in depth and zero trust security posture that uses multiple layers, including physical, technical, and administrative security controls to protect and defend your data.

Azure Key Vault and Azure Key Vault Managed HSM are designed, deployed, and operated so that *Microsoft and its agents are precluded from accessing, using, or extracting any data that's stored in the service, including cryptographic keys*.

Customer-managed keys that are securely created in or securely imported to the HSM devices, unless set otherwise by the customer, aren't extractable and are never visible in plaintext to Microsoft systems, employees, or agents.

The Key Vault team explicitly doesn't have operating procedures for granting this kind of access to Microsoft and its agents, even if authorized by a customer.

We won't attempt to defeat customer-controlled encryption features like Azure Key Vault or Azure Key Vault Managed HSM. If faced with a legal demand to do so, we would challenge such a demand on any lawful basis, consistent with our customer commitments.

Next, we take a detailed look at how these security controls are implemented.

Security controls in Azure Key Vault Managed HSM

Azure Key Vault Managed HSM uses the following types of security controls:

- Physical
- Technical
- Administrative

Physical security controls

The core of Managed HSM is the hardware security module (HSM). An HSM is a specialized, hardened, tamper-resistant, high-entropy, dedicated cryptographic processor that is validated to the FIPS 140-2 Level 3 standard. All components of the HSM are further covered in hardened epoxy and a metal casing to keep your keys safe from an attacker. The HSMs are housed in racks of servers across several datacenters, regions, and geographies. These geographically dispersed datacenters comply with key industry standards such as ISO/IEC 27001:2013 and NIST SP 800-53 for security and reliability.

Microsoft designs, builds, and operates datacenters in a way that strictly controls physical access to the areas where your keys and data are stored. Additional layers of

physical security, such as tall fences made of concrete and steel, dead-bolt steel doors, thermal alarm systems, closed-circuit live camera monitoring, 24x7 security personnel, need-to-access basis with approval per floor, rigorous staff training, biometrics, background checks, and access request and approval are mandated. The HSM devices and related servers are locked in a cage, and cameras film the front and back of the servers.

Technical security controls

Several layers of technical controls in Managed HSM further protect your key material. But most importantly, they prevent Microsoft from accessing the key material.

- **Confidentiality**: The Managed HSM service runs inside a trusted execution environment that's built on Intel Software Guard Extensions (Intel SGX). Intel SGX offers enhanced protection from internal and external attackers by using hardware isolation in enclaves that protect data in use.

 Enclaves are secured portions of the hardware's processor and memory. You can't view data or code inside the enclave, even with a debugger. If untrusted code tries to change content in the enclave memory, Intel SGX disables the environment and denies the operation.

 These unique capabilities help you protect your cryptographic key material from being accessible or visible as clear text. Also, Azure confidential computing offers solutions to enable the isolation of your sensitive data while it's being processed in the cloud.

- **Security domain**: A security domain is an encrypted blob that contains extremely sensitive cryptographic information. The security domain contains artifacts like the HSM backup, user credentials, the signing key, and the data encryption key that's unique to your managed HSM.

 The security domain is generated both in the managed HSM hardware and in the service software enclaves during initialization. After the managed HSM is provisioned, you must create at least three RSA key pairs. You send the public keys to the service when you request the security domain download. After the security domain is downloaded, the managed HSM moves into an activated state and is ready for you to consume. Microsoft personnel have no way of recovering the security domain, and they can't access your keys without the security domain.

- **Access controls and authorization**: Access to a managed HSM is controlled through two interfaces, the management plane and the data plane.

The management plane is where you manage the HSM itself. Operations in this plane include creating and deleting managed HSMs and retrieving managed HSM properties.

The data plane is where you work with the data that's stored in a managed HSM, which is HSM-backed encryption keys. From the data plane interface, you can add, delete, modify, and use keys to perform cryptographic operations, manage role assignments to control access to the keys, create a full HSM backup, restore a full backup, and manage security domain.

To access a managed HSM in either plane, all callers must have proper authentication and authorization. *Authentication* establishes the identity of the caller. *Authorization* determines which operations the caller can execute. A caller can be any one of the security principals that are defined in Azure Active Directory: User, group, service principal, or managed identity.

Both planes use Azure Active Directory for authentication. For authorization, they use different systems:

- The management plane uses Azure role-based access control (Azure RBAC), an authorization system that's built on Azure Resource Manager.
- The data plane uses a managed HSM-level RBAC (Managed HSM local RBAC), an authorization system that's implemented and enforced at the managed HSM level. The local RBAC control model allows designated HSM administrators to have complete control over their HSM pool that even the management group, subscription, or resource group administrators can't override.
- **Encryption in transit**: All traffic to and from the Managed HSM is always encrypted with TLS (Transport Layer Security versions 1.3 and 1.2 are supported) to protect against data tampering and eavesdropping where the TLS termination happens inside the SGX enclave and not in the untrusted host
- **Firewalls**: Managed HSM can be configured to restrict who can reach the service in the first place, which further shrinks the attack surface. We allow you to configure Managed HSM to deny access from the public internet and only allow traffic from trusted Azure services (such as Azure Storage)
- **Private endpoints**: By enabling a private endpoint, you're bringing the Managed HSM service into your virtual network allowing you to isolate that service only to trusted endpoints like your virtual network and Azure services. All traffic to and from your managed HSM will travel along the secure Microsoft backbone network without having to traverse the public internet.
- **Monitoring and logging**: The outermost layer of protection is the monitoring and logging capabilities of Managed HSM. By using the Azure Monitor service, you can check your logs for analytics and alerts to ensure that access patterns conform with your expectations. This allows members of your security team to

have visibility into what is happening within the Managed HSM service. If something doesn't look right, you can always roll your keys or revoke permissions.

- **Bring your own key (BYOK)**: BYOK enables Azure customers to use any supported on-premises HSMs to generate keys, and then import them to the managed HSM. Some customers prefer to use on-premises HSMs to generate keys to meet regulatory and compliance requirements. Then, they use BYOK to securely transfer an HSM-protected key to the managed HSM. The key to be transferred never exists outside of an HSM in plaintext form. During the import process, the key material is protected with a key that's held in the managed HSM.
- **External HSM**: Some customers have asked us if they can explore the option of having the HSM outside the Azure cloud to keep the data and keys segregated with an external HSM, either on a third-party cloud or on-premises. Although using a third-party HSM outside of Azure seems to give customers more control over keys, it introduces several concerns, such as latency that causes performance issues, SLA slip that's caused by issues with the third-party HSM, and maintenance and training costs. Also, a third-party HSM can't use key Azure features like soft delete and purge protection. We continue to evaluate this technical option with our customers to help them navigate the complex security and compliance landscape.

Administrative security controls

These administrative security controls are in place in Azure Key Vault Managed HSM:

- **Data defense**. You have Microsoft's strong commitment to challenge government requests and to defend your data.
- **Contractual obligations**. It offers control obligations for security and customer data protection as discussed in Microsoft Trust Center.
- **Cross-region replication**. Soon, you can use geo replication in Managed HSM to deploy HSMs in a secondary region.
- **Disaster recovery**. Azure offers an end-to-end backup and disaster recovery solution that is simple, secure, scalable, and cost-effective:
 - Business continuity management program
 - Azure Site Recovery
 - Azure Backup
 - Azure well-architected framework
- **Microsoft Security Response Center (MSRC)**. Managed HSM service administration is tightly integrated with MSRC.

 - Security monitoring for unexpected administrative operations with full 24/7 security response
- **Cloud resilient and secure supply chain**. Managed HSM advances reliability through a resilient cloud supply chain.
- **Regulatory compliance built-in initiative**. Compliance in Azure Policy provides built-in initiative definitions to view a list of the controls and compliance domains based on responsibility (Customer, Microsoft, Shared). For Microsoft-responsible controls, we provide additional details of our audit results based on third-party attestation and our implementation details to achieve that compliance.
- **Audit reports**. Resources to help information security and compliance professionals understand cloud features, and to verify technical compliance and control requirements
- **Assume breach philosophy**. We assume that any component could be compromised at any time, and we design and test appropriately. We do regular Red Team/Blue Team exercises (attack simulation).

Managed HSM offers robust physical, technical, and administrative security controls. Managed HSM gives you sole control of your key material for a scalable, centralized cloud key management solution that helps satisfy growing compliance, security, and privacy needs. Most importantly it provides encryption safeguards that are required for compliance. Our customers can be assured that we're committed to ensuring that their data is protected with transparency about our practices as we progress toward the implementation of the Microsoft EU Data Boundary.

For more information, reach out to your Azure account team to facilitate a discussion with the Azure Key Management product team.

Encryption in the Microsoft Cloud

Article • 09/08/2023

Customer data within Microsoft's enterprise cloud services is protected by several technologies and processes, including various forms of encryption. (Customer data in this document includes Exchange Online mailbox content, e-mail body, calendar entries, and the content of e-mail attachments, and if applicable, Skype for Business content, SharePoint Online site content and the files stored within sites, and files uploaded to OneDrive for Business or Skype for Business.) Microsoft uses multiple encryption methods, protocols, and ciphers across its products and services to help provide a secure path for customer data to travel through our cloud services, and to help protect the confidentiality of customer data that is stored within our cloud services. Microsoft uses some of the strongest, most secure encryption protocols available to provide barriers against unauthorized access to customer data. Proper key management is also an essential element of encryption best practices, and Microsoft works to ensure that all Microsoft-managed encryption keys are properly secured.

Customer data stored within Microsoft's enterprise cloud services is protected using one or more forms of encryption. (Validation of our crypto policy and its enforcement is independently verified by multiple third-party auditors, and reports of those audits are available on the Service Trust Portal .)

Microsoft provides service-side technologies that encrypt customer data at rest and in transit. For example, for customer data at rest, Microsoft Azure uses BitLocker and DM-Crypt , and Microsoft 365 uses BitLocker, Azure Storage Service Encryption, Distributed Key Manager (DKM), and Microsoft 365 service encryption. For customer data in transit, Azure, Office 365, Microsoft Commercial Support, Microsoft Dynamics 365, Microsoft Power BI, and Visual Studio Team Services use industry-standard secure transport protocols, such as Internet Protocol Security (IPsec) and Transport Layer Security (TLS), between Microsoft datacenters and between user devices and Microsoft datacenters.

In addition to the baseline level of cryptographic security provided by Microsoft, our cloud services also include cryptography options that you can manage. For example, you can enable encryption for traffic between their Azure virtual machines (VMs) and their users. With Azure Virtual Networks , you can use the industry-standard IPsec protocol to encrypt traffic between your corporate VPN gateway and Azure. You can also encrypt traffic between the VMs on your virtual network. In addition, new Office 365 Message Encryption capabilities allow you to send encrypted mail to anyone.

Following the Public Key Infrastructure Operational Security Standard, which is a component of the Microsoft Security Policy , Microsoft uses the cryptographic

capabilities included in the Windows operating system for certificates and authentication mechanisms. These mechanisms include the use of cryptographic modules that meet the U.S. government's Federal Information Processing Standards (FIPS) 140-2 standard. You can search for the relevant NIST certificate numbers for Microsoft using the Cryptographic Module Validation Program CMVP.

> [NOTE] To access the Microsoft Security Policy as a resource, you must sign in using your work or school account. If you don't have a subscription yet, you can sign up for a free trial.

FIPS 140-2 is a standard designed specifically for validating product modules that implement cryptography rather than the products that use them. Cryptographic modules that are implemented within a service can be certified as meeting the requirements for hash strength, key management, and the like. The cryptographic modules and ciphers used to protect the confidentiality, integrity, or availability of data in Microsoft's cloud services meet the FIPS 140-2 standard.

Microsoft certifies the underlying cryptographic modules used in our cloud services with each new release of the Windows operating system:

- Azure and Azure U.S. Government
- Dynamics 365 and Dynamics 365 U.S. Government
- Office 365, Office 365 U.S. Government, and Office 365 U.S. Government Defense

Encryption of customer data at rest is provided by multiple service-side technologies, including BitLocker, DKM, Azure Storage Service Encryption, and service encryption in Exchange Online, Skype for Business, OneDrive for Business, and SharePoint Online. Microsoft 365 service encryption includes an option to use customer-managed encryption keys that are stored in Azure Key Vault. This customer-managed key option, called Customer Key, is available for Exchange Online, SharePoint Online, Skype for Business, OneDrive for Business, Teams files, and Windows 365 Cloud PCs (in public preview).

For customer data in transit, all Office 365 servers negotiate secure sessions using TLS by default with client machines to secure customer data. For example, Office 365 will negotiate secure sessions to Skype for Business, Outlook, and Outlook on the web, mobile clients, and web browsers.

(All customer-facing servers negotiate to TLS 1.2 by default.)

Tip

If you're not an E5 customer, use the 90-day Microsoft Purview solutions trial to explore how additional Purview capabilities can help your organization manage data security and compliance needs. Start now at the **Microsoft Purview compliance portal trials hub**. Learn details about **signing up and trial terms**.

Related Links

- Encryption in Azure
- BitLocker and Distributed Key Manager (DKM) for Encryption
- Service encryption and key management
- Microsoft 365 encryption for Skype for Business, OneDrive for Business, SharePoint Online, and Exchange Online
- Encryption for data in transit
- Customer-managed encryption features
- Encryption risks and protections
- Encryption in Microsoft Dynamics 365

Feedback

Was this page helpful? Yes No

BitLocker and Distributed Key Manager (DKM) for Encryption

Article • 09/08/2023

Microsoft servers use BitLocker to encrypt the disk drives containing customer data at rest at the volume-level. BitLocker encryption is a data protection feature that is built into Windows. BitLocker is one of the technologies used to safeguard against threats in case there are lapses in other processes or controls (e.g., access control or recycling of hardware) that could lead to someone gaining physical access to disks containing customer data. In this case, BitLocker eliminates the potential for data theft or exposure because of lost, stolen, or inappropriately decommissioned computers and disks.

BitLocker is deployed with Advanced Encryption Standard (AES) 256-bit encryption on disks containing customer data in Exchange Online, SharePoint Online, and Skype for Business. Disk sectors are encrypted with a Full Volume Encryption Key (FVEK), which is encrypted with the Volume Master Key (VMK), which in turn is bound to the Trusted Platform Module (TPM) in the server. The VMK directly protects the FVEK and therefore, protecting the VMK becomes critical. The following figure illustrates an example of the BitLocker key protection chain for a given server (in this case, using an Exchange Online server).

The following table describes the BitLocker key protection chain for a given server (in this case, an Exchange Online server).

KEY PROTECTOR	GRANULARITY	HOW GENERATED?	WHERE IS IT STORED?	PROTECTION
AES 256-bit External Key	Per Server	BitLocker APIs	TPM or Secret Safe	Lockbox / Access Control
			Mailbox Server Registry	TPM encrypted
48-digit Numerical Password	Per Disk	BitLocker APIs	Active Directory	Lockbox / Access Control
X.509 Certificate as Data Recovery Agent (DRA) also	Environment (for example, Exchange Online multitenant)	Microsoft CA	Build System	No one user has the full password to the private key. The

KEY PROTECTOR	GRANULARITY	HOW GENERATED?	WHERE IS IT STORED?	PROTECTION
called Public Key Protector				password is under physical protection.

BitLocker key management involves the management of recovery keys that are used to unlock/recover encrypted disks in a Microsoft data center. Microsoft 365 stores the master keys in a secured share, only accessible by individuals who have been screened and approved. The credentials for the keys are stored in a secured repository for access control data (what we call a "secret store"), which requires a high level of elevation and management approvals to access using a just-in-time access elevation tool.

BitLocker supports keys which fall into two management categories:

- BitLocker-managed keys, which are short-lived and tied to the lifetime of an operating system instance installed on a server or to a given disk. These keys are deleted and reset during server reinstallation or disk formatting.
- BitLocker recovery keys, which are managed outside of BitLocker but used for disk decryption. BitLocker uses recovery keys for the scenario in which an operating system is reinstalled, and encrypted data disks already exist. Recovery keys are also used by Managed Availability monitoring probes in Exchange Online where a responder may need to unlock a disk.

BitLocker-protected volumes are encrypted with a full volume encryption key, which in turn is encrypted with a volume master key. BitLocker uses FIPS-compliant algorithms to ensure that encryption keys are never stored or sent over the wire in the clear. The Microsoft 365 implementation of customer data-at-rest-protection doesn't deviate from the default BitLocker implementation.

How Exchange Online uses TLS to secure email connections

Article • 11/27/2023

Learn how Exchange Online and Microsoft 365 use Transport Layer Security (TLS) and Forward Secrecy (FS) to secure email communications.

> **Tip**
>
> If you're not an E5 customer, use the 90-day Microsoft Purview solutions trial to explore how additional Purview capabilities can help your organization manage data security and compliance needs. Start now at the **Microsoft Purview compliance portal trials hub**. Learn details about **signing up and trial terms**.

TLS basics for Microsoft 365 and Exchange Online

Transport Layer Security (TLS), and Secure Sockets Layer (SSL) that came before TLS, are cryptographic protocols. These protocols secure communication over a network by using security certificates to encrypt a connection between computers. TLS supersedes SSL and is often referred to as SSL 3.1. Exchange Online uses TLS to encrypt the connections between Exchange servers and the connections between Exchange servers and other servers. For example, TLS is used to encrypt the connection between Exchange Online and your on-premises Exchange servers or your recipients' mail servers. Once the connection is encrypted, all data sent through that connection is sent through the encrypted channel.

TLS doesn't encrypt the message, just the connection. So, if you forward a message that was sent through a TLS-encrypted connection to a recipient organization that doesn't support TLS encryption, that message isn't necessarily encrypted.

If you want to encrypt the message, use an encryption technology that encrypts the message contents. For example, you can use Microsoft Purview Message Encryption or S/MIME. See Email encryption in Office 365 and Message encryption for information on message encryption in Office 365.

Use TLS in situations where you want to set up a secure channel of correspondence between Microsoft and your on-premises organization or another organization, such as a partner. Exchange Online always attempts to use TLS first to secure your email but

can't if the other party doesn't offer TLS security. Keep reading to find out how you can secure all mail to your on-premises servers or important partners by using *connectors*.

To provide the best-in-class encryption to our customers, Microsoft deprecated Transport Layer Security (TLS) versions 1.0 and 1.1 in Office 365 and Office 365 GCC. However, you can continue to use an unencrypted SMTP connection without any TLS. We don't recommend email transmission without any encryption.

How Exchange Online uses TLS between Exchange Online customers

Exchange Online servers always encrypt connections to other Exchange Online servers in our data centers with TLS 1.2. When you send a message to a recipient that is within your organization, Exchange Online automatically sends the message over an encrypted connection using TLS. Exchange Online also sends email that you send to other customers over encrypted connections using TLS that are secured using Forward Secrecy.

How Microsoft 365 uses TLS between Microsoft 365 and external, trusted partners

By default, Exchange Online always uses *opportunistic TLS*. Opportunistic TLS means Exchange Online always tries to encrypt connections with the most secure version of TLS first, then works its way down the list of TLS ciphers until it finds one on which both parties can agree. Unless you configure Exchange Online to ensure that messages to that recipient must use a secure connection, then by default Exchange sends the message without encryption if the recipient's organization doesn't support TLS encryption. Opportunistic TLS is sufficient for most businesses. However, for businesses that have compliance requirements such as medical, banking, or government organizations, you can configure Exchange Online to require, or force, TLS. For instructions, see Configure mail flow using connectors in Office 365.

If you decide to configure TLS between your organization and a trusted partner organization, Exchange Online can use *forced TLS* to create trusted channels of communication. Forced TLS requires your partner organization to authenticate to Exchange Online with a security certificate to send mail to you. Your partner needs to manage their own certificates. Exchange Online uses connectors to protect messages that you send from unauthorized access before they arrive at the recipient's email provider. For information on using connectors to configure mail flow, see Configure mail flow using connectors in Office 365.

TLS and hybrid Exchange Server deployments

If you're managing a hybrid Exchange deployment, your on-premises Exchange server needs to authenticate to Microsoft 365 using a security certificate to send mail to recipients whose mailboxes are only in Office 365. As a result, you need to manage your own security certificates for your on-premises Exchange servers. You must also securely store and maintain these server certificates. For more information about managing certificates in hybrid deployments, see Certificate requirements for hybrid deployments.

How to set up forced TLS for Exchange Online in Office 365

For Exchange Online customers, in order for forced TLS to work to secure all of your sent and received email, you need to set up more than one connector that requires TLS. You need one connector for messages sent to user mailboxes and another connector for messages sent from user mailboxes. Create these connectors in the Exchange admin center in Office 365. For instructions, see Configure mail flow using connectors in Office 365.

TLS certificate information for Exchange Online

The certificate information used by Exchange Online is described in the following table. If your business partner is setting up forced TLS on their email server, you need to provide this information to them. For security reasons, our certificates do change from time to time. The current certificate is valid from September 24, 2020.

Current certificate information valid from September 24, 2020

Attribute	Value
Certificate authority root issuer	DigiCert CA - 1
Certificate name	mail.protection.outlook.com
Organization	Microsoft Corporation
Organization unit	www.digicert.com
Certificate key strength	2048

Get more information about TLS, certificates, and Microsoft 365 and download certificates

Microsoft 365 encryption chains and certificate downloads

Microsoft 365 encryption chains and certificate downloads - DOD and GCC High

For a list of supported cipher suites, see Technical reference details about encryption.

Set up connectors for secure mail flow with a partner organization

Connectors with enhanced email security

Encryption in Microsoft 365

Feedback

Was this page helpful? Yes No

How Exchange Online secures your email secrets

Article • 09/11/2023

This article describes how Microsoft secures your email secrets in its data centers.

> **Tip**
>
> If you're not an E5 customer, use the 90-day Microsoft Purview solutions trial to explore how additional Purview capabilities can help your organization manage data security and compliance needs. Start now at the **Microsoft Purview compliance portal trials hub**. Learn details about **signing up and trial terms**.

How we secure secret information provided by you

In addition to the Office 365 Trust Center that provides Security, Privacy, and Compliance Information for Office 365, we use a technology called Distributed Key Manager (DKM).

Distributed Key Manager (DKM) is a client-side technology that uses a set of secret keys to encrypt and decrypt information. Only members of a specific security group in Active Directory Domain Services can access those keys in order to decrypt the data that is encrypted by DKM. In Exchange Online, only certain service accounts under which the Exchange processes run are part of that security group. No human is given credentials that are part of this security group and therefore no human has access to the keys that can decrypt these secrets.

For debugging, troubleshooting, or auditing purposes, a data center administrator must request elevated access to gain temporary credentials that are part of the security group. This process requires multiple levels of legal approval. If access is granted, all activity is logged and audited. Access is only granted for a set interval of time after which it automatically expires.

For extra protection, DKM technology includes automated key rollover and archiving. Automated rollover and archiving ensure that you can continue to access your older content without having to rely on the same key indefinitely.

Where Exchange Online uses DKM

Microsoft uses Distributed Key Manager to encrypt your secrets in Exchange Online data centers. For example:

- Email account credentials for connected accounts. Connected accounts are third-party accounts such as Hotmail, Gmail, and Yahoo! mail accounts.
- Customer Key. If you're using Customer Key, you'll use Azure Key Vault to safeguard your secrets.

Related articles

Encryption in Office 365

Technical reference details about encryption

Service assurance in the Microsoft Purview compliance portal

Feedback

Was this page helpful? Yes No

What is Microsoft 365 service encryption?

Article • 12/19/2023

In addition to using volume-level encryption, Exchange Online, Microsoft Teams, SharePoint Online, OneDrive for Business, and Windows 365 Cloud PCs also use service encryption to encrypt customer data.

Service encryption allows for two key management options, Microsoft-managed keys and customer-managed keys.

Service encryption provides multiple benefits:

- Provides a layer of protection for all Microsoft 365 services and Windows 365 Cloud PCs. For Microsoft 365 services, service encryption is an extra layer of protection is on top of BitLocker.
- Provides separation of Windows operating system administrators from access to application data stored or processed by the operating system.
- Includes a Customer Key option that enables multitenant services to provide per-tenant key management.
- Enhances the ability of Microsoft 365 and Windows 365 to meet your specific compliance requirements regarding encryption.

> **Tip**
>
> If you're not an E5 customer, use the 90-day Microsoft Purview solutions trial to explore how additional Purview capabilities can help your organization manage data security and compliance needs. Start now at the **Microsoft Purview compliance portal trials hub**. Learn details about **signing up and trial terms**.

What are Microsoft-managed keys?

By default, Microsoft manages all cryptographic keys including the root keys for service encryption. This option, called *Microsoft-managed keys*, is enabled by default for Exchange Online, SharePoint Online, OneDrive for Business, and Windows 365 Cloud PCs. Microsoft-managed keys provide default service encryption unless you decide to onboard using Customer Key. If, at a later date, you decide to stop using Customer Key

without following the data purge path, then your data stays encrypted using the Microsoft-managed keys. Your data is always encrypted at this default level at a minimum.

What is Customer Key?

Customer Key is a customer-managed key solution. With Customer Key, you supply root keys used with service encryption and you manage these keys using Azure Key Vault while Microsoft manages all other keys. Customer Key is available for Exchange Online, SharePoint Online, OneDrive for Business, and Windows 365 Cloud PCs.

Windows 365 support for Microsoft Purview Customer Key is in public preview and is subject to change.

Using Customer Key, you generate your own cryptographic keys. You can use an on-premises Hardware Service Module (HSM) or Azure Key Vault (AKV) to generate your keys. AKV lets you control and manage the cryptographic keys used by Microsoft 365. Customer Key uses the keys stored in the AKV as the root of one of the key chains that encrypts your mailbox data or files.

Customer Key provides you with more control over how Microsoft processes your data. For example, You can use Customer Key as a technical control if you want to terminate service with Microsoft or remove a portion of your data stored in the cloud. Removing data ensures no one, including Microsoft, can access or process the data. Customer Key is in addition and complementary to Customer Lockbox that you use to control access to your data by Microsoft personnel.

To learn how to set up Customer Key for Exchange Online, Microsoft Teams, SharePoint Online, including Team Sites, and OneDrive for Business, see these articles:

- Overview of Customer Key
- Set up Customer Key
- Manage Customer Key
- Roll or rotate a Customer Key or an availability key
- Understand the availability key

Feedback

Was this page helpful? Yes No

Email encryption

Article • 09/08/2023

This article compares encryption options in Microsoft 365 including Microsoft Purview Message Encryption, S/MIME, Information Rights Management (IRM), and introduces Transport Layer Security (TLS).

Microsoft 365 delivers multiple encryption options to help you meet your business needs for email security. This article presents three ways to encrypt email in Office 365. If you want to learn more about all security features in Office 365, visit the Office 365 Trust Center. This article introduces the three types of encryption available for Microsoft 365 administrators to help secure email in Office 365:

- Microsoft Purview Message Encryption.
- Secure/Multipurpose Internet Mail Extensions (S/MIME).
- Information Rights Management (IRM).

> **Tip**
>
> If you're not an E5 customer, use the 90-day Microsoft Purview solutions trial to explore how additional Purview capabilities can help your organization manage data security and compliance needs. Start now at the **Microsoft Purview compliance portal trials hub**. Learn details about **signing up and trial terms**.

How Microsoft 365 uses email encryption

Encryption is the process by which information is encoded so that only an authorized recipient can decode and consume the information. Microsoft 365 uses encryption in two ways: in the service, and as a customer control. In the service, encryption is used in Microsoft 365 by default; you don't have to configure anything. For example, Microsoft 365 uses Transport Layer Security (TLS) to encrypt the connection, or session, between two servers.

Here's how email encryption typically works:

- A message is encrypted, or transformed from plain text into unreadable ciphertext, either on the sender's machine, or by a central server while the message is in transit.

- The message remains in ciphertext while it's in transit in order to protect it from being read in case the message is intercepted.
- Once the message is received by the recipient, the message is transformed back into readable plain text in one of two ways:
 - The recipient's machine uses a key to decrypt the message, or
 - A central server decrypts the message on behalf of the recipient, after validating the recipient's identity.

For more information on how Microsoft 365 secures communication between servers, such as between organizations within Microsoft 365 or between Microsoft 365 and a trusted business partner outside of Microsoft 365, see How Exchange Online uses TLS to secure email connections in Office 365.

Comparing email encryption options available in Office 365

Email encryption technology			
What is it?	Message encryption is a service built on Azure Rights Management (Azure RMS) that lets you send encrypted email to people inside or outside your organization, regardless of the destination email address (Gmail, Yahoo! Mail, Outlook.com, etc.). As an admin, you can set up transport rules that define the conditions for encryption. When a user sends a message that matches a rule, encryption is applied automatically. To view encrypted messages, recipients can	IRM is an encryption solution that also applies usage restrictions to email messages. It helps prevent sensitive information from being printed, forwarded, or copied by unauthorized people. IRM capabilities in Microsoft 365 use Azure Rights Management (Azure RMS).	S/MIME is a certificate-based encryption solution that allows you to both encrypt and digitally sign a message. The message encryption helps ensure that only the intended recipient can open and read the message. A digital signature helps the recipient validate the identity of the sender. Both digital signatures and message encryption are made possible through the use of unique digital certificates that contain the keys for verifying digital signatures and

Email encryption technology	OME 	IRM 	S/MIME 
	either get a one-time passcode, sign in with a Microsoft account, or sign in with a work or school account associated with Office 365. Recipients can also send encrypted replies. They don't need a Microsoft 365 subscription to view encrypted messages or send encrypted replies.		encrypting or decrypting messages. To use S/MIME, you must have public keys on file for each recipient. Recipients have to maintain their own private keys, which must remain secure. If a recipient's private keys are compromised, the recipient needs to get a new private key and redistribute public keys to all potential senders.
What does it do?	OME: Encrypts messages sent to internal or external recipients. Allows users to send encrypted messages to any email address, including Outlook.com, Yahoo! Mail, and Gmail. Allows you, as an admin, to customize the email viewing portal to reflect your organization's brand. Microsoft securely manages and stores the keys, so you don't have to. No special client side software is needed as long as the encrypted message (sent as an HTML attachment) can be opened in a browser.	IRM: Uses encryption and usage restrictions to provide online and offline protection for email messages and attachments. Gives you, as an admin, the ability to set up transport rules or Outlook protection rules to automatically apply IRM to select messages. Lets users manually apply templates in Outlook or Outlook on the web (formerly known as Outlook Web App).	S/MIME addresses sender authentication with digital signatures, and message confidentiality with encryption.
What does it not do?	OME doesn't let you apply usage restrictions to messages. For	Some applications may not support IRM emails on all	S/MIME doesn't allow encrypted messages to

Email encryption technology	OME 	IRM 	S/MIME 
	example, you can't use it to stop a recipient from forwarding or printing an encrypted message.	devices. For more information about these and other products that support IRM email, see Client device capabilities.	be scanned for malware, spam, or policies.
Recommendations and example scenarios	We recommend using OME when you want to send sensitive business information to people outside your organization, whether they're consumers or other businesses. For example: A bank employee sending credit card statements to customers A doctor's office sending medical records to a patient An attorney sending confidential legal information to another attorney	We recommend using IRM when you want to apply usage restrictions as well as encryption. For example: A manager sending confidential details to her team about a new product applies the "Do Not Forward" option. An executive needs to share a bid proposal with another company, which includes an attachment from a partner who is using Office 365, and require both the email and the attachment to be protected.	We recommend using S/MIME when either your organization or the recipient's organization requires true peer-to-peer encryption. S/MIME is most commonly used in the following scenarios: Government agencies communicating with other government agencies A business communicating with a government agency

Don't apply multiple email encryption technologies to the same email message. Some email clients, such as Outlook for Mac, Outlook for iOS, and Outlook for Android, aren't able to open messages with multiple email encryption technologies applied.

What encryption options are available for my Microsoft 365 subscription?

For information about email encryption options for your Microsoft 365 subscription see the Exchange Online service description. Here, you can find information about the following encryption features:

- Azure RMS, including both IRM capabilities and Microsoft Purview Message Encryption
- S/MIME
- TLS
- Encryption of data at rest (through BitLocker, service encryption, and DKM)

You can also use third-party encryption tools with Microsoft 365, for example, PGP (Pretty Good Privacy). Microsoft 365 does not support PGP/MIME and you can only use PGP/Inline to send and receive PGP-encrypted emails.

What about encryption for data at rest?

"Data at rest" refers to data that isn't actively in transit. In Microsoft 365, email data at rest is encrypted using BitLocker Drive Encryption. BitLocker encrypts the hard drives in Microsoft data centers to provide enhanced protection against unauthorized access. To learn more, see BitLocker Overview.

More information about email encryption options

For more information about the email encryption options in this article as well as TLS, see these articles:

Microsoft Purview Message Encryption

Message encryption

IRM

Information Rights Management in Exchange Online

What is Azure Rights Management?

S/MIME

S/MIME for message signing and encryption

Understanding S/MIME

Understanding Public Key Cryptography

TLS

Configure custom mail flow by using connectors

How Exchange Online uses TLS to secure email connections

Disabling TLS 1.0 and 1.1 for Microsoft 365

Technical reference details about encryption

Article • 09/08/2023

Refer to this article for information about certificates, technologies, and TLS cipher suites used for encryption in Microsoft 365. This article also provides details about planned deprecations.

- If you're looking for overview information, see Encryption in Microsoft 365.
- If you're looking for setup information, see Set up encryption in Microsoft 365 Enterprise.
- For specific information about TLS 1.1 and 1.0 deprecation, see Disabling TLS 1.0 and 1.1 for Microsoft 365.
- For information about cipher suites supported by specific versions of Windows, see Cipher Suites in TLS/SSL (Schannel SSP).
- For certificate chains, see Microsoft 365 encryption chains and Microsoft 365 encryption chains - DOD and GCC High.

> **Tip**
>
> If you're not an E5 customer, use the 90-day Microsoft Purview solutions trial to explore how additional Purview capabilities can help your organization manage data security and compliance needs. Start now at the **Microsoft Purview compliance portal trials hub**. Learn details about **signing up and trial terms**.

Microsoft Office 365 certificate ownership and management

You don't need to purchase or maintain certificates for Office 365. Instead, Office 365 uses its own certificates.

Current encryption standards and planned deprecations

To provide best-in-class encryption, Office 365 regularly reviews supported encryption standards. Sometimes, old standards are deprecated as they become out of date and less secure. This article describes currently supported cipher suites and other standards and details about planned deprecations.

FIPS compliance for Microsoft 365

All cipher suites supported by Office 365 use algorithms acceptable under FIPS 140-2. Office 365 inherits FIPS validations from Windows (through Schannel). For information about Schannel, see Cipher Suites in TLS/SSL (Schannel SSP).

AES256-CBC support for Microsoft 365

In late August 2023, Microsoft Purview Information Protection will begin to use Advanced Encryption Standard (AES) with 256-bit key length in Cipher Block Chaining mode (AES256-CBC). By October 2023, AES256-CBC will be the default for encryption of Microsoft 365 Apps documents and emails. You may need to take action to support this change in your organization.

Who is impacted and what do I need to do?

Use this table to figure out if you have to take action:

Client applications	Service applications	Action required?	What do I need to do?
Microsoft 365 Apps	Exchange Online, SharePoint Online	No	N/A
Office 2013, 2016, 2019, or 2021	Exchange Online, SharePoint Online	Yes (Optional)	See Set up Office 2013, 2016, 2019, or 2021 for AES256-CBC mode.
Microsoft 365 Apps	Exchange Server or hybrid	Yes (Mandatory)	See Set up Exchange Server for AES256-CBC support.
Office 2013, 2016, 2019, or 2021	Exchange Server or hybrid	Yes (Mandatory)	Complete Option 1 (required), and then see Set up Office 2013, 2016, 2019, or 2021 for AES256-CBC mode.
Microsoft 365 Apps	MIP SDK	Yes (Optional)	See Set up MIP SDK for AES256-CBC support.
Any	SharePoint Server	No	N/A

Set up Office 2013, 2016, 2019, or 2021 for AES256-CBC mode

You need to configure Office 2013, 2016, 2019, or 2021 to use AES256-CBC mode using Group Policy, or by using the Cloud Policy service for Microsoft 365. Starting with version 16.0.16327 of Microsoft 365 Apps, CBC mode is used by default. Use the `Encryption mode for Information Rights Management (IRM)` setting under `User Configuration/Administrative Templates/Microsoft Office 2016/Security Settings`.

For example, to force CBC mode, select the group policy setting as follows:

Encryption mode for Information Rights Management (IRM): [1, Cipher Block Chaining (CBC)]

Set up Exchange Server for AES256-CBC support

Exchange Server doesn't support decrypting content that uses AES256-CBC. To work around this problem, you have two options.

Option 1

Customers using Exchange Online with the Azure Rights Management Connector service deployed will be opted out of the AES256-CBC publishing change in both Exchange Online and SharePoint Online.

To move to AES256-CBC mode, complete these steps:

1. Install the hotfix on your Exchange Servers when it becomes available. For the most recent information about ship dates, see the Microsoft 365 product roadmap.
2. If you're using Exchange Server with the Azure Rights Management Connector Service, you'll need to run the GenConnectorConfig.ps1 script on each Exchange server. For more information, see Configure servers for the Rights Management connector. To download the Azure RMS connector, see the official Microsoft Download Center.

Once your organization has installed the patch across all of your Exchange Servers, open a support case and request these services to be enabled for AES256-CBC publishing.

Option 2

This option gives you some extra time before you need to patch all of your Exchange servers. Use this option if you're unable to complete the steps in option 1 when the hotfix becomes available. Instead, deploy group policy or client settings that force Microsoft 365 clients to keep using AES128-ECB mode. Deploy this setting using Group Policy, or by using the Cloud Policy service for Microsoft 365. You can configure Office and Microsoft 365 Apps for Windows to use ECB or CBC mode with the `Encryption mode for Information Rights Management (IRM)` setting under `User Configuration/Administrative Templates/Microsoft Office 2016/Security Settings`. Starting with version 16.0.16327 of Microsoft 365 Apps, CBC mode is used by default.

For example, to force EBC mode for Windows clients, set the group policy setting as follows:

Encryption mode for Information Rights Management (IRM): [2, Electronic Codebook (ECB)]

To configure settings for Office for Mac clients, see Set suite-wide preferences for Office for Mac.

As soon as you can, complete the steps in option 1.

Set up MIP SDK for AES256-CBC support

Update to MIP SDK 1.13 or later. If you choose to update to MIP SDK 1.13, you'll need to configure a setting to force AES256-CBC. For more information, see the MIP SDK Version 1.13.158 Critical Update. Later versions of the MIP SDK will protect Microsoft 365 files and email with AES256-CBC by default.

Versions of TLS supported by Microsoft 365

TLS, and SSL that came before TLS, are cryptographic protocols that secure communication over a network by using security certificates to encrypt a connection between computers. Microsoft 365 supports TLS version 1.2 (TLS 1.2).

Some of the services continue to support TLS version 1.3 (TLS 1.3).

> **Important**
>
> Be aware that TLS versions deprecate, and that deprecated versions *should not be used* where newer versions are available. If your legacy services do not require TLS 1.0 or 1.1 you should disable them.

Support for TLS 1.0 and 1.1 deprecation

Office 365 stopped supporting TLS 1.0 and 1.1 on October 31, 2018. We have completed disabling TLS 1.0 and 1.1 in GCC High and DoD environments. We began disabling TLS 1.0 and 1.1 for Worldwide and GCC environments beginning on October 15, 2020 and will continue with roll-out over the next weeks and months.

To maintain a secure connection to Office 365 and Microsoft 365 services, all client-server and browser-server combinations use TLS 1.2 and modern cipher suites. You might have to update certain client-server and browser-server combinations. For information about how this change impacts you, see Preparing for the mandatory use of TLS 1.2 in Office 365.

Deprecating support for 3DES

Since October 31, 2018, Microsoft 365 no longer supports the use of 3DES cipher suites for communication to Microsoft 365. More specifically, Microsoft 365 no longer supports the TLS_RSA_WITH_3DES_EDE_CBC_SHA cipher suite. Since February 28, 2019, this cipher suite has been disabled in Microsoft 365. Clients and servers that communicate with Microsoft 365 must support one or more of the supported ciphers. For a list of supported ciphers, see TLS cipher suites supported by Microsoft 365.

Deprecating SHA-1 certificate support in Microsoft 365

Since June 2016, Microsoft 365 no longer accepts an SHA-1 certificate for outbound or inbound connections. Use SHA-2 (Secure Hash Algorithm 2) or a stronger hashing algorithm in the certificate chain.

TLS cipher suites supported by Microsoft 365

TLS uses *cipher suites*, collections of encryption algorithms, to establish secure connections. Microsoft 365 supports the cipher suites listed in the following table. The table lists the cipher suites in order of strength, with the strongest cipher suite listed first.

Microsoft 365 responds to a connection request by first attempting to connect using the most secure cipher suite. If the connection doesn't work, Microsoft 365 tries the second most secure cipher suite in the list, and so on. The service continues down the list until the connection is accepted. Likewise, when Microsoft 365 requests a connection, the receiving service chooses whether to use TLS and which cipher suite to use.

Cipher suite name	Key exchange algorithm/strength	Forward secrecy	Cipher/strength	Authentication algorithm/strength
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	ECDH/192	Yes	AES/256	RSA/112
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	ECDH/128	Yes	AES/128	RSA/112
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	ECDH/192	Yes	AES/256	RSA/112
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	ECDH/128	Yes	AES/128	RSA/112
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	ECDH/192	Yes	AES/256	RSA/112
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	ECDH/128	Yes	AES/128	RSA/112
TLS_RSA_WITH_AES_256_GCM_SHA384	RSA/112	No	AES/256	RSA/112
TLS_RSA_WITH_AES_128_GCM_SHA256	RSA/112	No	AES/256	RSA/112

The following cipher suites supported TLS 1.0 and 1.1 protocols until their deprecation date. For GCC High and DoD environments that deprecation date was January 15, 2020. For Worldwide and GCC environments that date was October 15, 2020.

Protocols	Cipher suite name	Key exchange algorithm/strength	Forward secrecy	Cipher/strength	Authentication algorithm/strength
TLS 1.0, 1.1, 1.2	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	ECDH/192	Yes	AES/256	RSA/112
TLS 1.0, 1.1, 1.2	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	ECDH/128	Yes	AES/128	RSA/112
TLS 1.0, 1.1, 1.2	TLS_RSA_WITH_AES_256_CBC_SHA	RSA/112	No	AES/256	RSA/112

Protocols	Cipher suite name	Key exchange algorithm/strength	Forward secrecy	Cipher/strength	Authentication algorithm/strength
TLS 1.0, 1.1, 1.2	TLS_RSA_WITH_AES_128_CBC_SHA	RSA/112	No	AES/128	RSA/112
TLS 1.0, 1.1, 1.2	TLS_RSA_WITH_AES_256_CBC_SHA256	RSA/112	No	AES/256	RSA/112
TLS 1.0, 1.1, 1.2	TLS_RSA_WITH_AES_128_CBC_SHA256	RSA/112	No	AES/256	RSA/112

Certain Office 365 products (including Microsoft Teams) use Azure Front Door to terminate TLS connections and route network traffic efficiently. At least one of the cipher suites supported by Azure Front Door over TLS 1.2 must be enabled to successfully connect to these products. For Windows 10 and above, we recommend enabling one or both of the ECDHE cipher suites for better security. Windows 7, 8, and 8.1 aren't compatible with Azure Front Door's ECDHE cipher suites and the DHE cipher suites have been provided for compatibility with those operating systems.

Related articles

TLS Cipher Suites in Windows 10 v1903

Encryption in Office 365

Set up encryption in Office 365 Enterprise

Schannel implementation of TLS 1.0 in Windows security status update: November 24, 2015

TLS/SSL Cryptographic Enhancements (Windows IT Center)

Preparing for TLS 1.2 in Office 365 and Office 365 GCC

What are the current cipher suites supported by Azure Front Door?

Encryption for data-in-transit

Article • 04/04/2023

In addition to protecting customer data at rest, Microsoft uses encryption technologies to protect customer data in transit. Data is in transit:

- When a client machine communicates with a Microsoft server;
- When a Microsoft server communicates with another Microsoft server; and
- When a Microsoft server communicates with a non-Microsoft server (for example, Exchange Online delivering email to a third-party email server).

Inter-data center communications between Microsoft servers take place over TLS or IPsec, and all customer-facing servers negotiate a secure session using TLS with client machines (for example, Exchange Online uses TLS 1.2 with 256-bit cipher strength is used (FIPS 140-2 Level 2-validated). (See Technical reference details about encryption for a list of TLS cipher suites supported by Microsoft 365.) This applies to the protocols that are used by clients such as Outlook, Skype for Business, Microsoft Teams, and Outlook on the web (for example, HTTP, POP3, etc.).

The public certificates are issued by Microsoft IT SSL using SSLAdmin, an internal Microsoft tool to protect confidentiality of transmitted information. All certificates issued by Microsoft IT have a minimum of 2048 bits in length, and Webtrust compliance requires SSLAdmin to make sure that certificates are issued only to public IP addresses owned by Microsoft. Any IP addresses that fail to meet this criterion are routed through an exception process.

All implementation details such as the version of TLS being used, whether Forward Secrecy (FS) is enabled, the order of cipher suites, etc., are available publicly. One way to see these details is to use a third-party website, such as Qualys SSL Labs. Below are the links to automated test pages from Qualys that display information for the following services:

- Office 365 Portal
- Exchange Online
- SharePoint Online
- Skype for Business (SIP)
- Skype for Business (Web)
- Exchange Online Protection
- Microsoft Teams

For Exchange Online Protection, URLs vary by tenant names; however, all customers can test Microsoft 365 using **microsoft-com.mail.protection.outlook.com**.

Encryption for Skype for Business, OneDrive for Business, SharePoint Online, Microsoft Teams, and Exchange Online

Article • 03/02/2023

Microsoft 365 is a highly secure environment that offers extensive protection in multiple layers: physical data center security, network security, access security, application security, and data security.

Skype for Business

Skype for Business customer data may be stored at rest in the form of files or presentations that are uploaded by meeting participants. The Web Conferencing server encrypts customer data using AES with a 256-bit key. The encrypted customer data is stored on a file share. Each piece of customer data is encrypted using a different randomly generated 256-bit key. When a piece of customer data is shared in a conference, the Web Conferencing server instructs the conferencing clients to download the encrypted customer data via HTTPS. It sends the corresponding key to clients so that the customer data can be decrypted. The Web Conferencing server also authenticates conferencing clients before it allows the clients access to conference customer data. When joining a Web conference, each conferencing client establishes a SIP dialog with the conferencing focus component running inside the front-end server over TLS first. The conferencing focus passes to the conference client an authentication cookie generated by the Web Conferencing server. The conferencing client then connects to the Web Conferencing server presenting the authentication cookie to be authenticated by the server.

SharePoint Online and OneDrive for Business

All customer files in SharePoint Online are protected by unique, per-file keys that are always exclusive to a single tenant. The keys are either created and managed by the SharePoint Online service, or when Customer Key is used, created and managed by customers. When a file is uploaded, encryption is performed by SharePoint Online within the context of the upload request, before being sent to Azure storage. When a file is downloaded, SharePoint Online retrieves the encrypted customer data from Azure storage based on the unique document identifier and decrypts the customer data

before sending it to the user. Azure storage has no ability to decrypt, or even identify or understand the customer data. All encryption and decryption happen in the same systems that enforce tenant isolation, which are Microsoft Entra ID and SharePoint Online.

Several workloads in Microsoft 365 store data in SharePoint Online, including Microsoft Teams, which stores all files in SharePoint Online, and OneDrive for Business, which uses SharePoint Online for its storage. All customer data stored in SharePoint Online is encrypted (with one or more AES 256-bit keys) and distributed across the datacenter as follows. (Every step of this encryption process is FIPS 140-2 Level 2 validated. For additional information about FIPS 140-2 compliance, see FIPS 140-2 Compliance.)

- Each file is split into one or more chunks, depending on file size. Each chunk is encrypted using its own unique AES 256-bit key.
- When a file is updated, the update is handled in the same way: the change is split into one or more chunks, and each chunk is encrypted with a separate unique key.
- These chunks – files, pieces of files, and update deltas – are stored as blobs in Azure storage that are randomly distributed across multiple Azure storage accounts.
- The set of encryption keys for these chunks of customer data is itself encrypted.
 - The keys used to encrypt the blobs are stored in the SharePoint Online Content Database.
 - The Content Database is protected by database access controls and encryption at rest. Encryption is performed using Transparent Data Encryption (TDE) in Azure SQL Database. (Azure SQL Database is a general-purpose relational database service in Microsoft Azure that supports structures such as relational data, JSON, spatial, and XML.) These secrets are at the service level for SharePoint Online, not at the tenant level. These secrets (sometimes referred to as the master keys) are stored in a separate secure repository called the Key Store. TDE provides security at rest for both the active database and the database backups and transaction logs.
 - When customers provide the optional key, the customer key is stored in Azure Key Vault, and the service uses the key to encrypt a tenant key, which is used to encrypt a site key, which is then used to encrypt the file level keys. Essentially, a new key hierarchy is introduced when the customer provides a key.
- The map used to re-assemble the file is stored in the Content Database along with the encrypted keys, separately from the master key needed to decrypt them.

- Each Azure storage account has its own unique credentials per access type (read, write, enumerate, and delete). Each set of credentials is held in the secure Key Store and is regularly refreshed. As described above, there are three different types of stores, each with a distinct function:
 - Customer data is stored as encrypted blobs in Azure storage. The key to each chunk of customer data is encrypted and stored separately in the Content Database. The customer data itself holds no clue as to how it can be decrypted.
 - The Content Database is a SQL Server database. It holds the map required to locate and reassemble the customer data blobs held in Azure storage as well as the keys needed to encrypt those blobs. However, the set of keys is itself encrypted (as explained above) and held in a separate Key Store.
 - The Key Store is physically separate from the Content Database and Azure storage. It holds the credentials for each Azure storage container and the master key to the set of encrypted keys held in the Content Database.

Each of these three storage components – the Azure blob store, the Content Database, and the Key Store – is physically separate. The information held in any one of the components is unusable on its own. Without access to all three, it is impossible to retrieve the keys to the chunks, decrypt the keys to make them usable, associate the keys with their corresponding chunks, decrypt each chunk, or reconstruct a document from its constituent chunks.

BitLocker certificates, which protect the physical disk volumes on machines in the datacenter, are stored in a secure repository (the SharePoint Online secret store) that is protected by the Farm key.

The TDE keys that protect the per-blob keys are stored in two locations:

- The secure repository, which houses the BitLocker certificates and is protected by the Farm Key; and
- In a secure repository managed by Azure SQL Database.

The credentials used to access the Azure storage containers are also held in the SharePoint Online secret store and delegated to each SharePoint Online farm as needed. These credentials are Azure storage SAS signatures, with separate credentials used to read or write data, and with policy applied so that they auto-expire every 60 days. Different credentials are used to read or write data (not both) and SharePoint Online farms are not given permissions to enumerate.

> ⓘ **Note**

For Office 365 U.S. Government customers, data blobs are stored in Azure U.S. Government Storage. In addition, access to SharePoint Online keys in Office 365 U.S. Government is limited to Office 365 staff that have been specifically screened. Azure U.S. Government operations staff do not have access to the SharePoint Online key store that is used for encrypting data blobs.

For more information about data encryption in SharePoint Online and OneDrive for Business, see Data Encryption in OneDrive for Business and SharePoint Online.

List Items in SharePoint Online

List Items are smaller chunks of customer data that are created ad-hoc or that can live more dynamically within a site, such as rows in a user-created list, individual posts in a SharePoint Online blog, or entries within a SharePoint Online wiki page. List items are stored in the Content Database (Azure SQL Database) and protected with TDE.

Encryption of data in transit

In OneDrive for Business and SharePoint Online, there are two scenarios in which data enters and exits the datacenters.

- **Client communication with the server** - Communication to SharePoint Online and OneDrive for Business across the Internet uses TLS connections.
- **Data movement between datacenters** - The primary reason to move data between datacenters is for geo-replication to enable disaster recovery. For instance, SQL Server transaction logs and blob storage deltas travel along this pipe. While this data is already transmitted by using a private network, it is further protected with best-in-class encryption.

Exchange Online

Exchange Online uses BitLocker for all mailbox data, and the BitLocker configuration is described in BitLocker for Encryption. Service-level encryption encrypts all mailbox data at the mailbox level.

In addition to service-encryption, Microsoft 365 supports Customer Key, which is built on top of service-encryption. Customer Key is a Microsoft-managed key option for Exchange Online service encryption that is also on Microsoft's Roadmap. This method of encryption provides increased protection not afforded by BitLocker because it provides separation of server administrators and the cryptographic keys necessary for decryption

of data, and because the encryption is applied directly to the data (in contrast with BitLocker, which applies encryption at the logical disk volume) any customer data copied from an Exchange server remains encrypted.

The scope for Exchange Online service encryption is customer data that is stored at rest within Exchange Online. (Skype for Business stores nearly all user-generated content within the user's Exchange Online mailbox and therefore inherits the service encryption feature of Exchange Online.)

Microsoft Teams

Teams uses TLS and MTLS to encrypt instant messages. All server-to-server traffic requires MTLS, regardless of whether the traffic is confined to the internal network or crosses the internal network perimeter.

This table summarizes the protocols used by Teams.

Traffic type	Encrypted by
Server-to-server	MTLS
Client-to-server (ex. instant messaging and presence)	TLS
Media flows (ex. audio and video sharing of media)	TLS
Audio and video sharing of media	SRTP/TLS
Signaling	TLS

Media Encryption

Media traffic is encrypted using Secure RTP (SRTP), a profile of Real-Time Transport Protocol (RTP) that provides confidentiality, authentication, and replay attack protection to RTP traffic. SRTP uses a session key generated by using a secure random number generator and is exchanged using the signaling TLS channel. Client-to-client media traffic is negotiated through a client-to-server connection signaling but is encrypted using SRTP when going client-to-client directly.

Teams uses a credentials-based token for secure access to media relays over TURN. Media relays exchange the token over a TLS-secured channel.

FIPS

Teams uses FIPS (Federal Information Processing Standard) compliant algorithms for encryption key exchanges. For more information on the implementation of FIPS, please see Federal Information Processing Standard (FIPS) Publication 140-2.

Governance, risk, and compliance overview

Article • 04/04/2023

How does Microsoft provide effective security governance across the enterprise?

Microsoft understands that effective security policies must be implemented consistently across the enterprise to protect Microsoft information systems and customers. Security policies must also account for variations in business functions and information systems to be universally applicable. To meet these requirements, Microsoft implements a comprehensive security governance program as a part of the Microsoft Policy Framework. Security governance falls under the Microsoft Security Policy (MSP).

The MSP organizes Microsoft's security policies, standards, and requirements so they can be implemented across all Microsoft engineering groups and business units. Individual business units are responsible for specific implementations of Microsoft security policies. For example, Microsoft 365 documents its security implementations in the Microsoft 365 Information Security Policy and the related Microsoft 365 Control Framework. Azure and Dynamics 365 document their security implementations in the Standard Operating Procedures (SOPs) and the Azure Control Framework. These security implementations align with the goals and objectives of the MSP.

Microsoft's security governance program is informed by and aligns with various regulatory and compliance frameworks. Security requirements are constantly evolving to account for new technologies, regulatory and compliance requirements, and security threats. Because of these changes, Microsoft regularly updates our security policies and supporting documents to protect Microsoft systems and customers, meet our commitments, and maintain customer trust.

How do Microsoft online services implement the Microsoft Security Policy (MSP)?

Microsoft 365 documents security implementations in the Microsoft 365 Information Security Policy. This policy aligns with the Microsoft Security Policy and governs the Microsoft 365 information system, including all Microsoft 365 environments and all resources involved in the collection, processing, maintenance, use, sharing,

dissemination, and disposal of data. Similarly, Azure and Dynamics 365 use the Microsoft Security Policy to govern their information system.

The information systems include the following components governed by the Microsoft 365 Information Security Policy (for Microsoft 365) and the Microsoft Security Policy (for Azure and Dynamics 365):

- Infrastructure: The physical and hardware components of Azure, Dynamics 365, and Microsoft 365 systems (facilities, equipment, and networks)
- Software: The programs and operating software of Azure, Dynamics 365, and Microsoft 365 systems (systems, applications, and utilities)
- People: The personnel involved in the operation and use of Azure, Dynamics 365, and Microsoft 365 systems (developers, operators, users, and managers)
- Procedures: The programmed and manual procedures involved in the operation of Azure, Dynamics 365, and Microsoft 365 systems
- Data: The information generated, collected, and processed by Azure, Dynamics 365, and Microsoft 365 systems (transaction streams, files, databases, and tables)

The Microsoft 365 Information Security Policy is supplemented by the Microsoft 365 Control Framework. The Microsoft 365 Control Framework details the minimum-security requirements for all Microsoft 365 services and information system components. It also references the legal and corporate requirements behind each control. The framework includes control activity names, descriptions, and guidance to ensure effective control implementations by service teams. Microsoft 365 uses the control framework to track control implementations for internal and external reporting. Similarly, Azure and Dynamics 365 record control implementations in the Azure Control Framework.

How do online services limit and track exceptions to established policies and procedures?

All exceptions to the Control Frameworks must have legitimate business justification and be approved by an appropriate governance entity within each online services team. Depending on the scope of the exception and the potential risk it represents, approval for exceptions may need to be obtained from a corporate vice president or higher. Exceptions are managed in a tracking tool where they are reviewed and approved for continued relevance.

How does Microsoft assess and manage risk across the enterprise?

Risk management is the process of identifying, assessing, and responding to threats or events that can impact Company or customer objectives. Risk management at Microsoft is designed to anticipate new threats and provide ongoing security for our cloud systems and the customers who use them.

Microsoft's risk management align to the Enterprise Risk Management (ERM) framework. ERM enables the overall enterprise risk management process and works with management across the enterprise to identify and ensure accountability for Microsoft's most significant risks.



Microsoft ERM enables common risk management principles across the enterprise so business units can independently facilitate consistent and comparative risk assessments. This coordination gives Microsoft the ability to aggregate and report risk information in a consolidated manner for management. ERM provides business units in Microsoft with common methodologies, tools, and goals for the risk management process. Microsoft 365 and other engineering groups and business units use these tools to conduct individual risk assessments as part of their own risk management programs under the guidance of ERM.

How do Microsoft online services work with ERM?

Each online service follows ERM guidance to manage risks across Microsoft services. The program focuses on aligning the ERM framework with existing Microsoft engineering, service operations, and compliance processes, making the Risk Management program more effective and efficient. Each online service's risk management activities ultimately roll up into and inform the ERM process.

As part of risk assessment activities, each online service analyzes design and operating effectiveness of controls implemented as part of the Microsoft Controls Framework (Framework). The Framework is a rationalized set of controls that, when properly implemented along with supporting compliance activities, allows engineering teams to comply with key regulations and certifications.

How do online services keep security and compliance requirements updated?

Governance, Risk, and Compliance teams of each online service (GRC) work to maintain the Control Framework on an ongoing basis. Several scenarios may require the GRC team to update the control framework, including changes in relevant regulations or laws, emerging threats, penetration test results, security incidents, audit feedback, and new compliance requirements. When a framework change is required, the Trust team identifies key stakeholders responsible for approving and implementing the change to ensure it is feasible and will not cause unintended issues with Online services. Once the GRC team and relevant stakeholders agree on what the change requires, the workloads responsible for implementing the change set target completion dates and work to implement the change within their respective services. After implementation targets have been met, the Trust team updates the control framework with the new or updated controls.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to governance, risk, and compliance.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
ISO 27001/27002 Statement of Applicability Certificate	A.5: Information security policies A.18.1: Compliance with legal and contractual requirements A.18.2: Information security reviews	April 24, 2023
ISO 27017 Statement of Applicability Certificate	A.5: Information security policies A.18.1: Compliance with legal and contractual requirements A.18.2: Information security reviews	April 24, 2023
ISO 27018 Statement of Applicability Certificate	A.5: Information security policies	April 24, 2023
ISO 22301 Certificate	6.1.1: Determining risks and opportunities 6.1.2: Addressing risks and opportunities	April 24, 2023
SOC 1	IS-1: Microsoft security policy IS-2: Microsoft security policy review IS-3: Security roles and responsibilities	November 17, 2023
SOC 2 SOC 3	C5-1: Standard operating procedures IS-1: Microsoft security policy IS-2: Microsoft security policy review IS-3: Security roles and responsibilities SOC2-14: Confidentiality and non-disclosure agreements SOC2-18: Statutory, regulatory, and contractual requirements SOC2-19: Cross-functional compliance program SOC2-20: ISMS program SOC2-26: Annual risk assessment	November 17, 2023

Microsoft 365

Expand table

External audits	Section	Latest report date
FedRAMP (Office 365)	CA-2: Security assessments CA-5: Plan of action and milestones PL-2: System security plan RA-3: Risk assessment	July 31, 2023
ISO 27001/27002/27017 Statement of Applicability Certification (27001/27002) Certification (27017)	A.5: Information security policies A.18.1: Compliance with legal and contractual requirements A.18.2: Information security reviews	March 2023
SOC 1	CA-03: Risk management	January 3, 2023
SOC 2	CA-02: Governance, risk, and compliance team responsibilities CA-03: Risk management CA-11: Policy framework updates CA-17: Microsoft security policy CA-24: Internal risk assessment CA-25: Control framework updates	January 3, 2023

Resources

- Microsoft Security Policy
- Microsoft Security Program Policy

Feedback

Was this page helpful? Yes No

Microsoft 365 Risk Management program

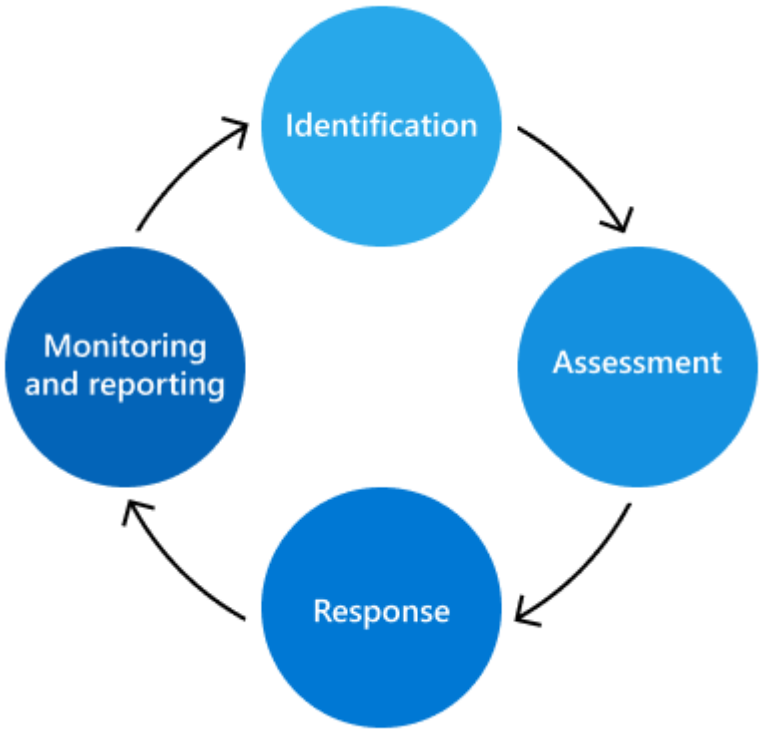
Article • 03/02/2023

The purpose of the Microsoft 365 Risk Management program is to identify, assess, and manage risks to Microsoft 365. Microsoft's top priority is to proactively identify and address risks that could impact our service infrastructure, as well as our customers, their data, and their trust. In addition, a robust risk management program is necessary to meet contractual obligations and maintain public accreditations that our customers rely on to satisfy their own compliance requirements. While the Microsoft 365 Risk Management program functions independently, it aligns with the overarching Enterprise Risk Management (ERM) program's policies, priorities, and methodologies. Working with the ERM program allows for consistent comparison across business units and engineering groups, contributing to a more cohesive approach to risk management across the enterprise.

The Microsoft 365 Trust team is responsible for managing the Microsoft 365 Risk Management program and conducting the activities defined by the ERM program. The Trust team focuses on integrating the risk management framework with existing Microsoft 365 engineering, service operations, and compliance process to make the Risk Management program more effective and efficient.

The Trust team also maintains the Microsoft 365 Controls Framework, a set of rationalized controls that, when properly implemented with supporting compliance activities, allows engineering teams to comply with key regulations and certifications. This framework is continuously updated based on feedback and findings as part of the risk management process.

Risk management activities fall into four phases: identification, assessment, response, and monitoring and reporting.



Identification

The risk management process starts with identifying all possible risks to all key control areas, internal and external threats, and vulnerabilities in the Microsoft 365 environment. The information guiding this process comes from multiple sources including interviews, vulnerability scans, attack simulation exercises, audit findings, and incident management activities.

The Trust team interviews subject matter experts (SMEs) from multiple service teams on previously identified risks and potential future risks that may be introduced as the services grow. Additionally, SMEs help to validate the accuracy and completeness of risks identified from the other continuous monitoring sources.

The identification phase is also when decision logs, active security and compliance exceptions, and mitigation work from previous risk assessments are reviewed.

Assessment

Each identified risk is assessed using three metrics: impact, likelihood, and control deficiency.

- Impact refers to the damage that would occur to the service, business, or Microsoft if that risk were to be realized. The impact to Microsoft may include damage to reputation, loss of customers, or legal/compliance implications.

- Likelihood defines the probability of the potential risk being realized and is calculated by analyzing the probability and frequency with which it will occur.
- Control deficiency measures the effectiveness of implemented mitigation controls.

These metrics are used to calculate a risk score that represents the severity of each risk, accounting for existing mitigation strategies. Risks are aggregated and presented to key stakeholders from each service to verify the accuracy and completeness of Microsoft 365's risk posture.

Response

Using the verified list of risks to Microsoft 365, the Trust team assigns risks to the affected service for risk response. Defined guidelines help determine the appropriate risk response strategy based on the risk score and control effectiveness. Risk response strategies fall into four categories:

- Tolerate: Areas of low-risk exposure with a low level of control.
- Operate: Areas of low-risk exposure where controls are deemed adequate.
- Monitor: Areas of high-risk exposure where controls are deemed adequate and should be monitored for effectiveness.
- Improve: Areas of high-risk exposure with a low level of control that are top priorities in addressing.

The Trust team coordinates with service teams to develop plans for addressing each risk. The severity level determines the appropriate level of review and approval for each plan. For risks that require action, existing engineering bug processes are used for tracking, managing, and making exception decisions. Using a process familiar to the engineering and operation teams makes risk response more efficient and effective.

Monitoring and reporting

Risks identified as part of the risk assessment are monitored and reported to relevant stakeholders. Monitoring strategies include security monitoring, periodic risk reviews, penetration testing, and vulnerability scanning. These monitoring efforts act as data sources for reporting on key performance indicators, creating dashboards, and developing formal reports, all of which inform future risk decisions.

Multiple times a year, the Trust team meets with risk owners from each service to review risk scores, evaluate the effectiveness of their action plans, and make updates where needed. In addition, Microsoft 365's risk assessment activities contribute to the ERM

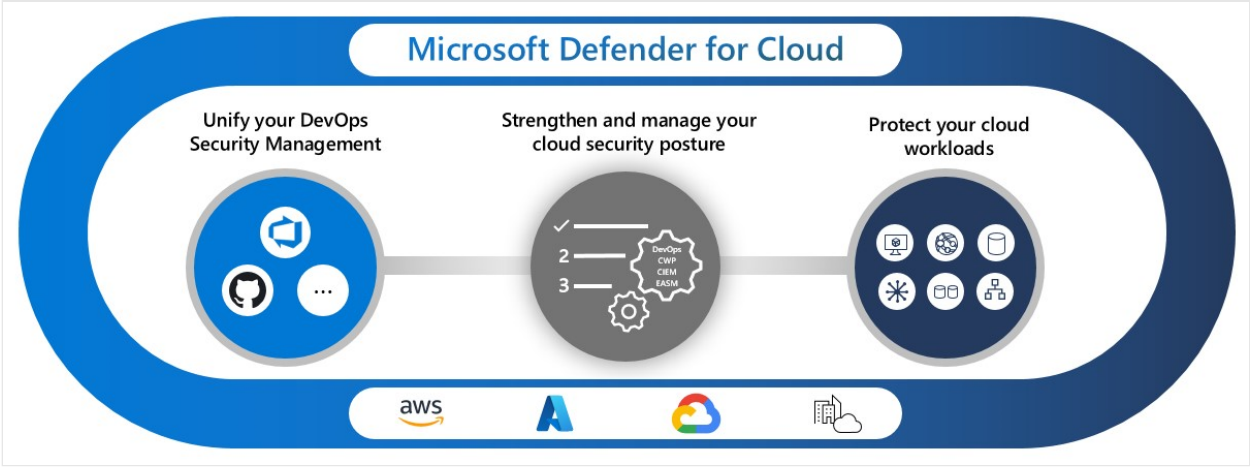
program's Enterprise Risk Assessments, which provide a high-level overview of Microsoft's risk posture to Microsoft senior management and the ERM program.

What is Microsoft Defender for Cloud?

Article • 11/16/2023

Microsoft Defender for Cloud is a cloud-native application protection platform (CNAPP) that is made up of security measures and practices that are designed to protect cloud-based applications from various cyber threats and vulnerabilities. Defender for Cloud combines the capabilities of:

- A development security operations (DevSecOps) solution that unifies security management at the code level across multicloud and multiple-pipeline environments
- A cloud security posture management (CSPM) solution that surfaces actions that you can take to prevent breaches
- A cloud workload protection platform (CWPP) with specific protections for servers, containers, storage, databases, and other workloads



> **Note**
>
> For Defender for Cloud pricing information, see the **pricing page**.

When you enable Defender for Cloud, you automatically gain access to Microsoft 365 Defender.

The Microsoft 365 Defender portal helps security teams investigate attacks across cloud resources, devices, and identities. Microsoft 365 Defender provides an overview of attacks, including suspicious and malicious events that occur in cloud environments. Microsoft 365 Defender accomplishes this goal by correlating all alerts and incidents, including cloud alerts and incidents.

You can learn more about the integration between Microsoft Defender for Cloud and Microsoft Defender XDR.

Secure cloud applications

Defender for Cloud helps you to incorporate good security practices early during the software development process, or DevSecOps. You can protect your code management environments and your code pipelines, and get insights into your development environment security posture from a single location. Defender for Cloud empowers security teams to manage DevOps security across multi-pipeline environments.

Today's applications require security awareness at the code, infrastructure, and runtime levels to make sure that deployed applications are hardened against attacks.

Expand table

Capability	What problem does it solve?	Get started	Defender plan
Code pipeline insights	Empowers security teams with the ability to protect applications and resources from code to cloud across multi-pipeline environments, including GitHub, Azure DevOps, and GitLab. DevOps security findings, such as Infrastructure as Code (IaC) misconfigurations and exposed secrets, can then be correlated with other contextual cloud security insights to prioritize remediation in code.	Connect Azure DevOps, GitHub, and GitLab repositories to Defender for Cloud	Foundational CSPM (Free) and Defender CSPM

Improve your security posture

The security of your cloud and on-premises resources depends on proper configuration and deployment. Defender for Cloud recommendations identifies the steps that you can take to secure your environment.

Defender for Cloud includes Foundational CSPM capabilities for free. You can also enable advanced CSPM capabilities by enabling the Defender CSPM plan.

Expand table

Capability	What problem does it solve?	Get started	Defender plan
Centralized policy	Define the security conditions that you want to maintain across your environment.	Customize a security policy	Foundational CSPM (Free)

Capability	What problem does it solve?	Get started	Defender plan
management	The policy translates to recommendations that identify resource configurations that violate your security policy. The Microsoft cloud security benchmark is a built-in standard that applies security principles with detailed technical implementation guidance for Azure and other cloud providers (such as AWS and GCP).		
Secure score	Summarize your security posture based on the security recommendations. As you remediate recommendations, your secure score improves.	Track your secure score	Foundational CSPM (Free)
Multicloud coverage	Connect to your multicloud environments with agentless methods for CSPM insight and CWP protection.	Connect your Amazon AWS and Google GCP cloud resources to Defender for Cloud	Foundational CSPM (Free)
Cloud Security Posture Management (CSPM)	Use the dashboard to see weaknesses in your security posture.	Enable CSPM tools	Foundational CSPM (Free)
Advanced Cloud Security Posture Management	Get advanced tools to identify weaknesses in your security posture, including: - Governance to drive actions to improve your security posture - Regulatory compliance to verify compliance with security standards - Cloud security explorer to build a comprehensive view of your environment	Enable CSPM tools	Defender CSPM
Data-aware Security Posture	Data-aware security posture automatically discovers datastores containing sensitive data, and helps reduce risk of data breaches.	Enable data-aware security posture	Defender CSPM or Defender for Storage
Attack path analysis	Model traffic on your network to identify potential risks before you implement changes to your environment.	Build queries to analyze paths	Defender CSPM
Cloud Security Explorer	A map of your cloud environment that lets you build queries to find security risks.	Build queries to find security risks	Defender CSPM
Security governance	Drive security improvements through your organization by assigning tasks to resource	Define governance	Defender CSPM

Capability	What problem does it solve?	Get started	Defender plan
	owners and tracking progress in aligning your security state with your security policy.	rules	
Microsoft Entra Permissions Management	Provide comprehensive visibility and control over permissions for any identity and any resource in Azure, AWS, and GCP.	Review your Permission Creep Index (CPI)	Defender CSPM

Protect cloud workloads

Proactive security principles require that you implement security practices that protect your workloads from threats. Cloud workload protections (CWP) surface workload-specific recommendations that lead you to the right security controls to protect your workloads.

When your environment is threatened, security alerts right away indicate the nature and severity of the threat so you can plan your response. After you identify a threat in your environment, you need to quickly respond to limit the risk to your resources.

Expand table

Capability	What problem does it solve?	Get started	Defender plan
Protect cloud servers	Provide server protections through Microsoft Defender for Endpoint or extended protection with just-in-time network access, file integrity monitoring, vulnerability assessment, and more.	Secure your multicloud and on-premises servers	Defender for Servers
Identify threats to your storage resources	Detect unusual and potentially harmful attempts to access or exploit your storage accounts using advanced threat detection capabilities and Microsoft Threat Intelligence data to provide contextual security alerts.	Protect your cloud storage resources	Defender for Storage
Protect cloud databases	Protect your entire database estate with attack detection and threat response for the most popular database types in Azure to protect the database engines and data types, according to their attack surface and security risks.	Deploy specialized protections for cloud and on-premises databases	- Defender for Azure SQL Databases - Defender for SQL servers on machines - Defender for Open-source relational

Capability	What problem does it solve?	Get started	Defender plan
			databases - Defender for Azure Cosmos DB
Protect containers	Secure your containers so you can improve, monitor, and maintain the security of your clusters, containers, and their applications with environment hardening, vulnerability assessments, and run-time protection.	Find security risks in your containers	Defender for Containers
Infrastructure service insights	Diagnose weaknesses in your application infrastructure that can leave your environment susceptible to attack.	- Identify attacks targeting applications running over App Service - Detect attempts to exploit Key Vault accounts - Get alerted on suspicious Resource Manager operations - Expose anomalous DNS activities	- Defender for App Service - Defender for Key Vault - Defender for Resource Manager - Defender for DNS
Security alerts	Get informed of real-time events that threaten the security of your environment. Alerts are categorized and assigned severity levels to indicate proper responses.	Manage security alerts	Any workload protection Defender plan
Security incidents	Correlate alerts to identify attack patterns and integrate with Security Information and Event Management (SIEM), Security Orchestration Automated Response (SOAR), and IT Service Management (ITSM) solutions to respond to threats and limit the risk to your resources.	Export alerts to SIEM, SOAR, or ITSM systems	Any workload protection Defender plan

ⓘ Important

As of August 1, customers with an existing subscription to Defender for DNS can continue to use the service, but new subscribers will receive alerts about suspicious DNS activity as part of Defender for Servers P2.

Learn More

For more information about Defender for Cloud and how it works, check out:

- A step-by-step walkthrough of Defender for Cloud
- An interview about Defender for Cloud with an expert in cybersecurity in Lessons Learned from the Field
- Microsoft Defender for Cloud - Use cases
- Microsoft Defender for Cloud PoC Series - Microsoft Defender for Containers

Next steps

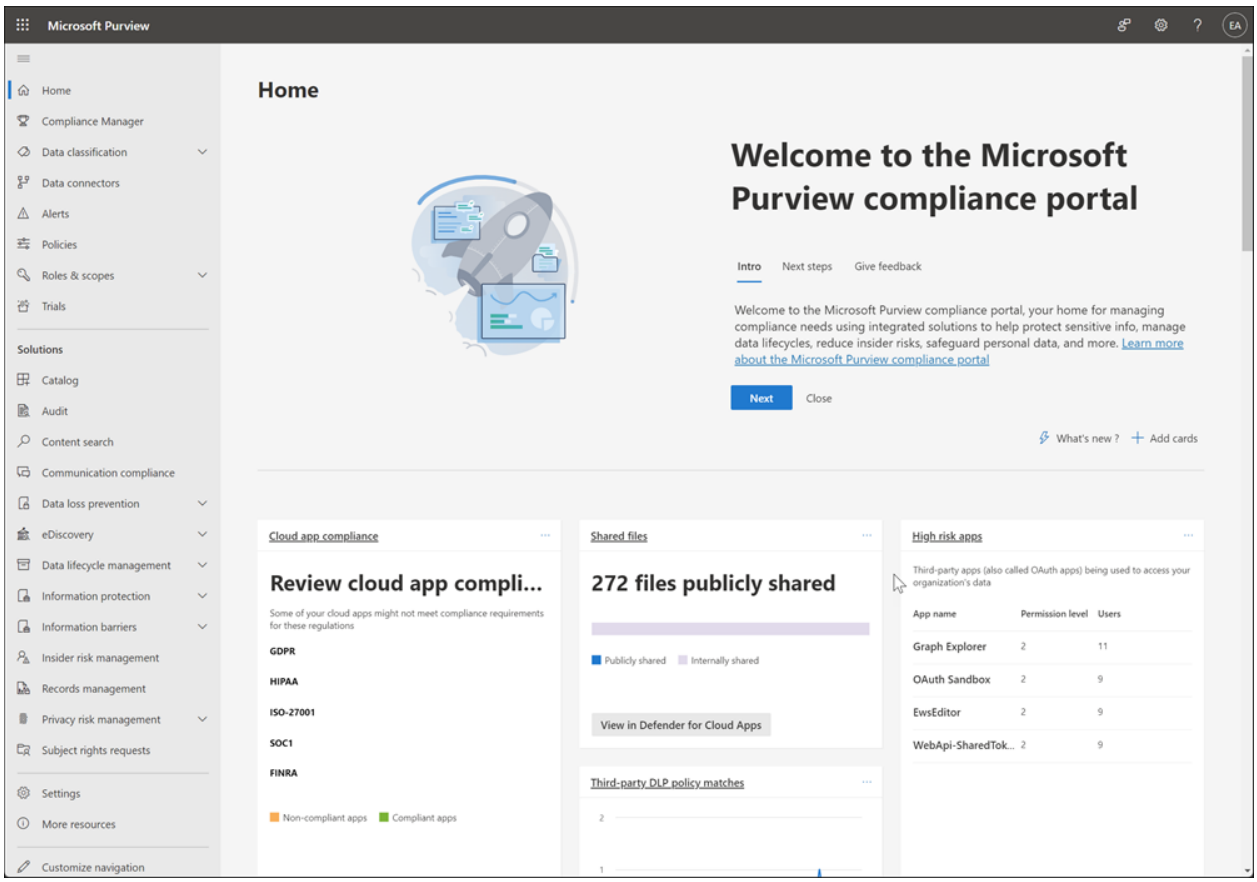
Enable Microsoft Defender plans

Microsoft Purview compliance portal

Article • 11/14/2023

The Microsoft Purview compliance portal is your platform for accessing all the risk and compliance solutions in Microsoft Purview. The portal provides access to the data and tools for managing your organization's compliance needs.

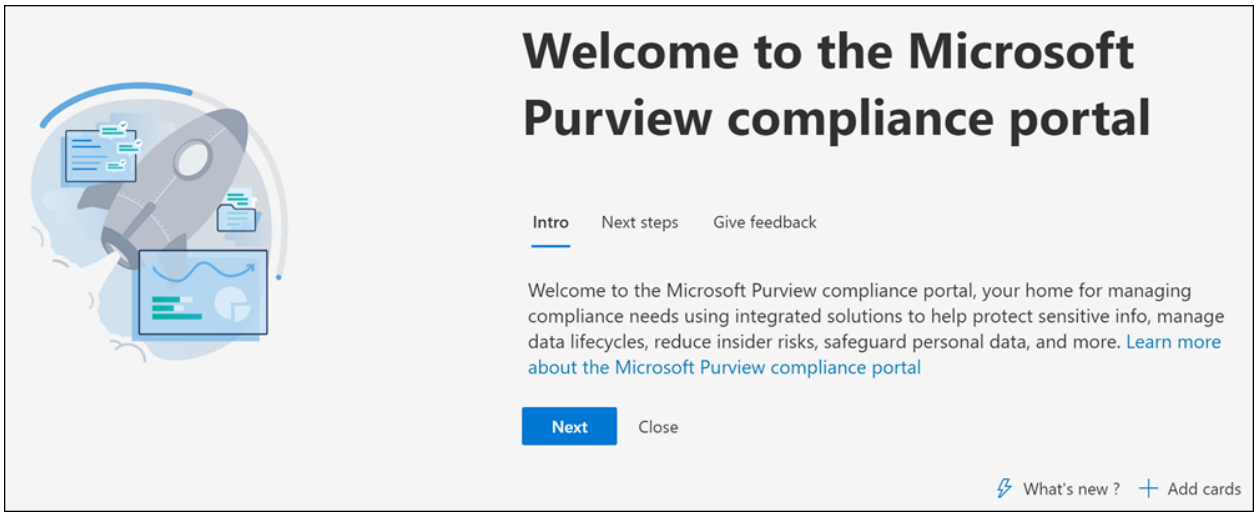
Read this article to get acquainted with the compliance portal, how to access it, and your next steps.



> **Tip**
>
> If you're not an E5 customer, use the 90-day Microsoft Purview solutions trial to explore how additional Purview capabilities can help your organization manage data security and compliance needs. Start now at the **Microsoft Purview compliance portal trials hub**. Learn details about **signing up and trial terms**.

Welcome to Microsoft Purview

When you go to the compliance portal for the first time, you're greeted with the following welcome message:



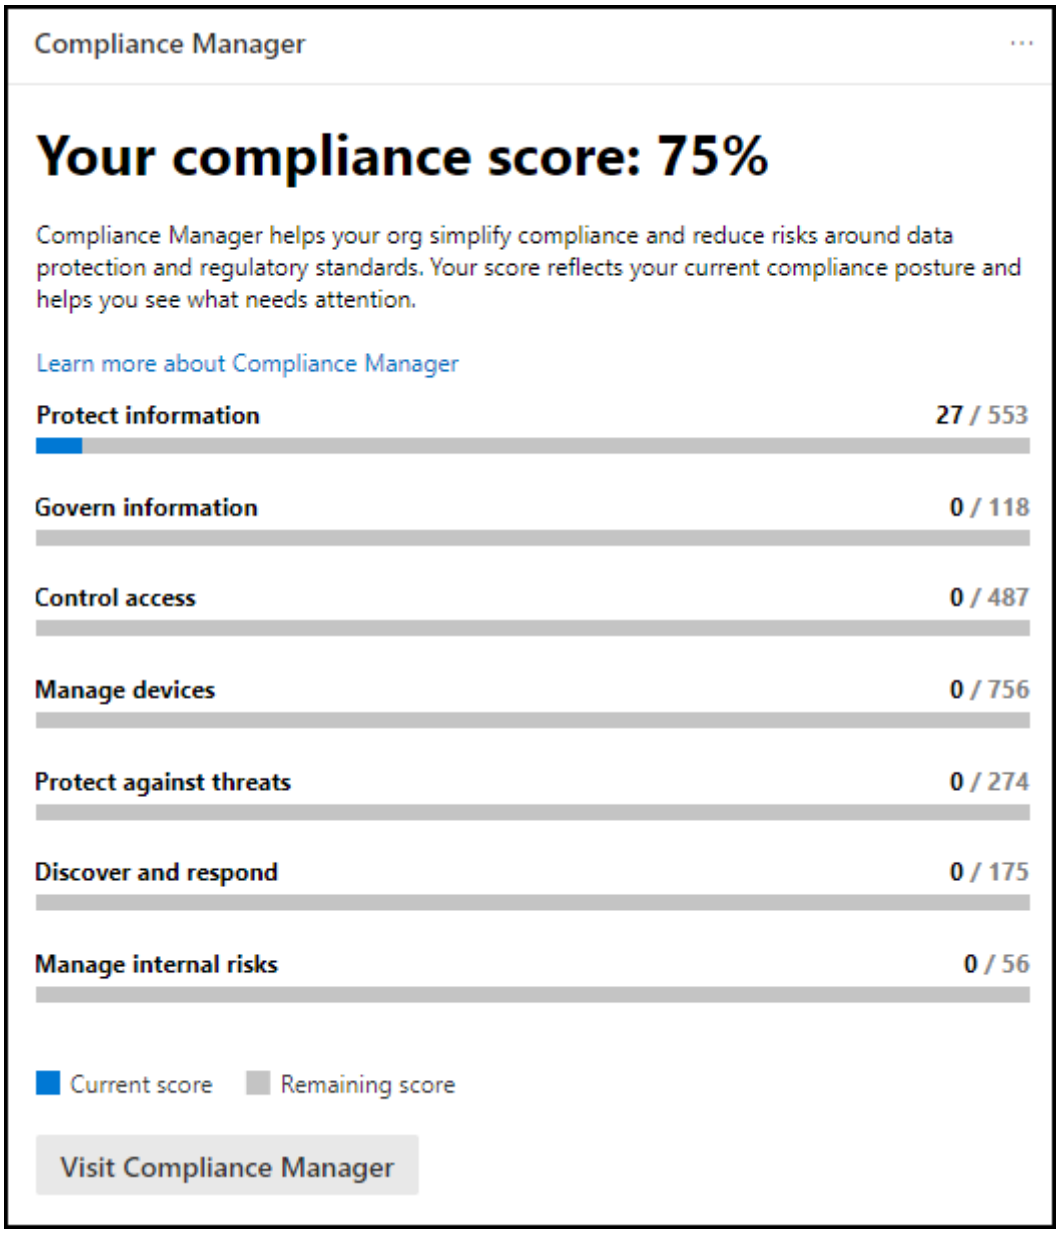
The welcome banner gives you some pointers on how to get started, with next steps, and an invitation for you to give us feedback.

Card section

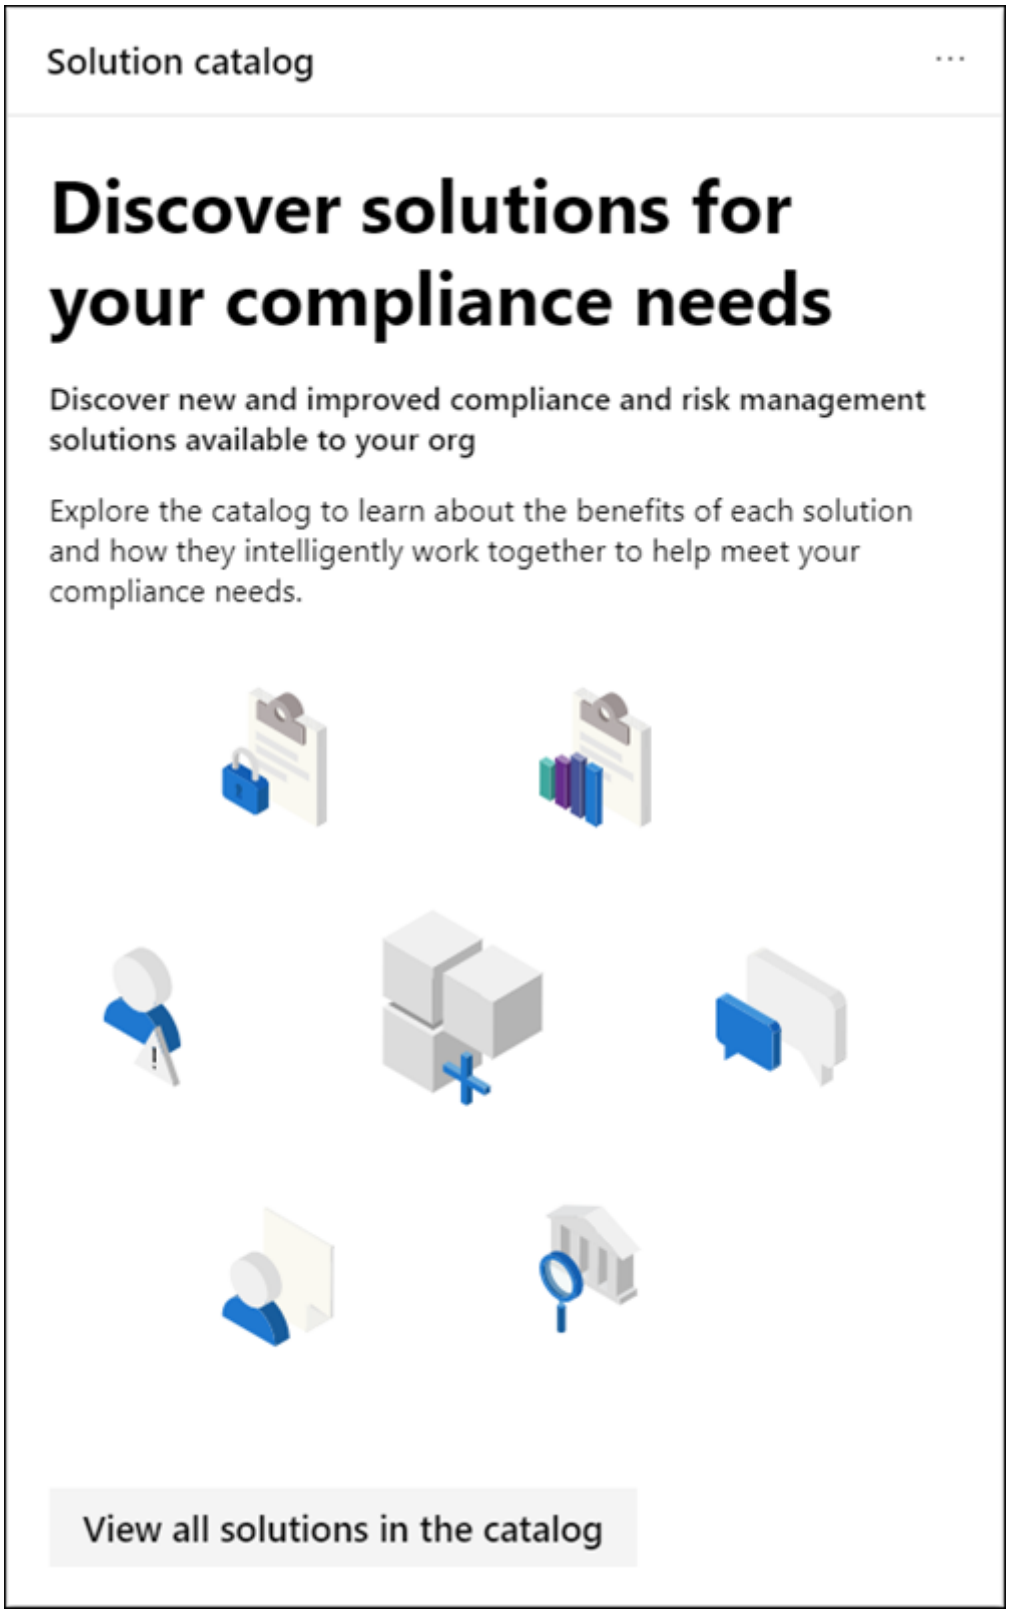
When you first visit the compliance portal, the card section on the home page shows you at a glance how your organization is doing with data compliance, what solutions are available for your organization, and a summary of any active alerts.

From here, you can:

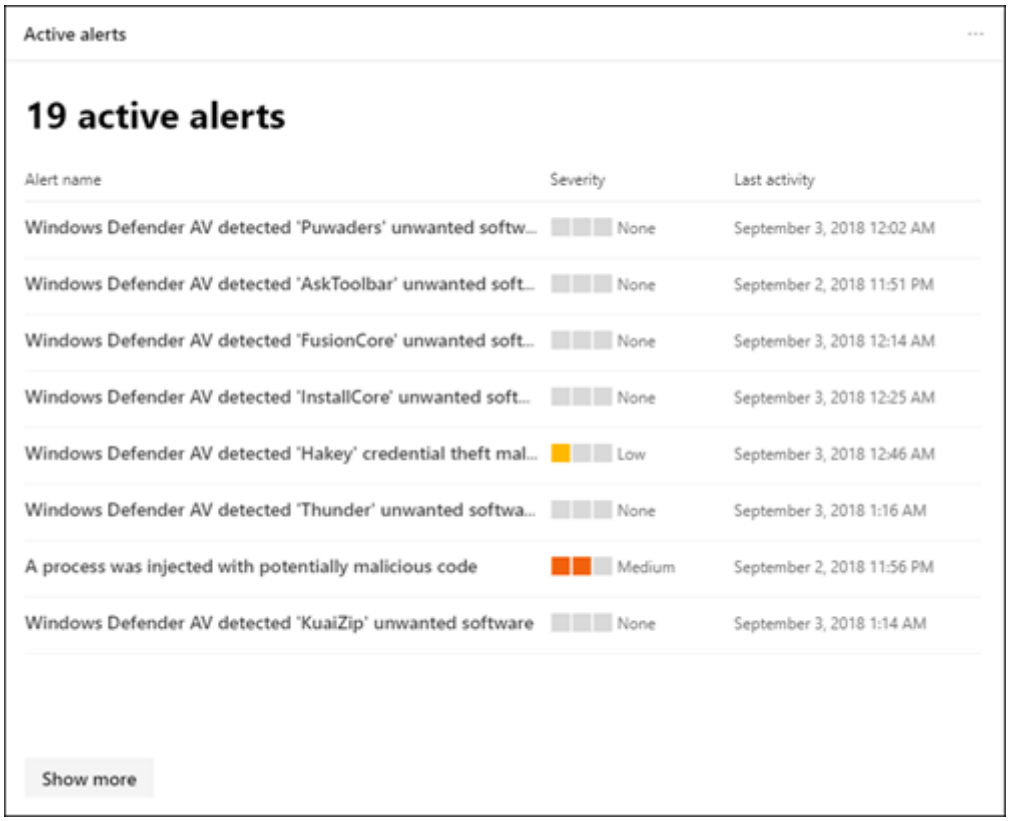
- Review the **Microsoft Purview Compliance Manager** card, which leads you to the Compliance Manager solution. Compliance Manager helps simplify the way you manage compliance. It calculates a risk-based score measuring your progress toward completing recommended actions that help reduce risks around data protection and regulatory standards. It also provides workflow capabilities and built-in control mapping to help you efficiently carry out improvement actions.



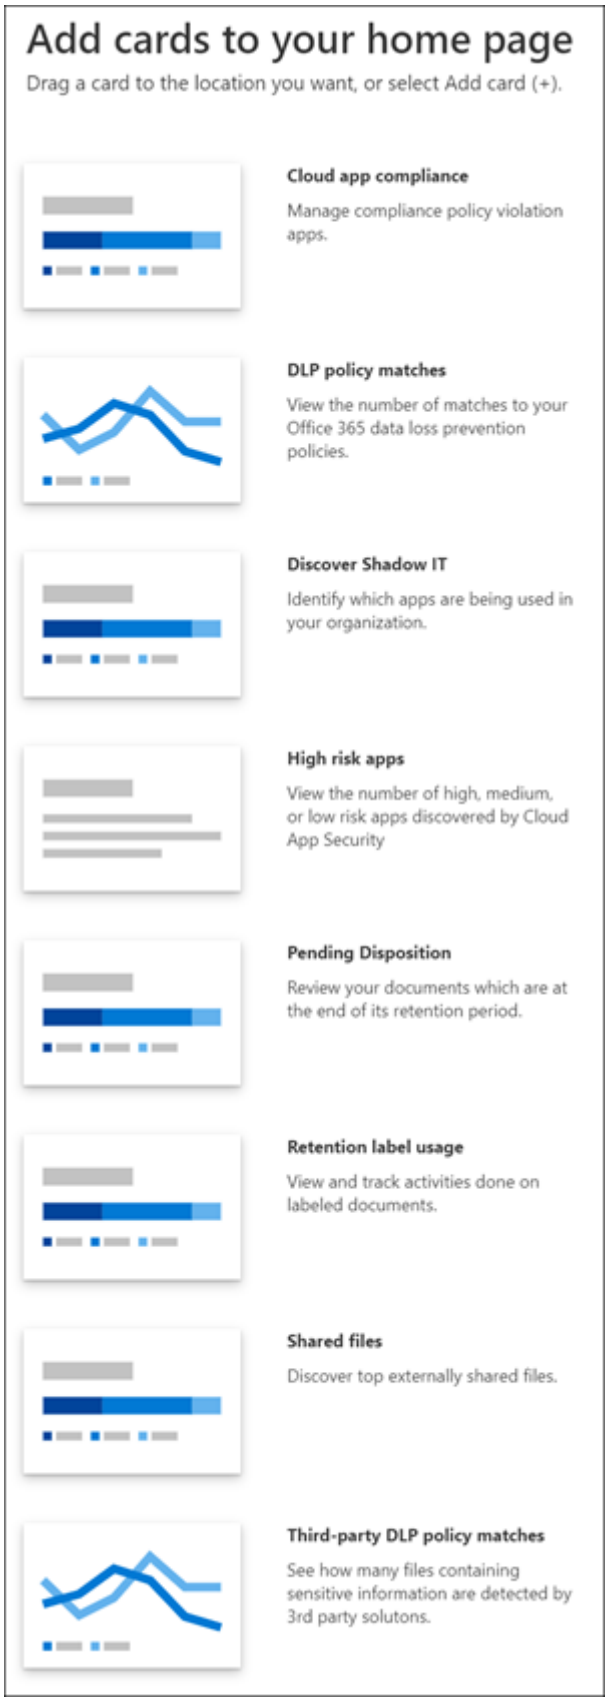
- Review the new **Solution catalog** card, which links to collections of integrated solutions you can use to help you manage end-to-end compliance scenarios. A solution's capabilities and tools might include a combination of policies, alerts, reports, and more.



- Review the **Active alerts** card, which includes a summary of the most active alerts and includes a link where you can view more detailed information, such as Severity, Status, Category, and more.



You can also use the **Add cards** feature to add additional cards, such as one showing your organization's cloud app compliance, and another showing data about users with shared files, with links to Defender for Cloud Apps or other tools where you can explore data.

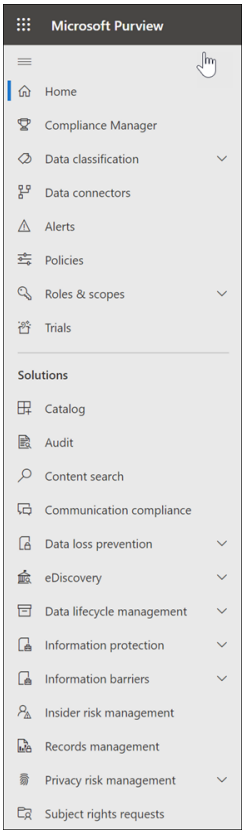


Easy navigation to more compliance features and capabilities

In addition to links in cards on the home page, you'll see a navigation pane on the left side of the screen that gives you easy access to your alerts, reports, policies, compliance

solutions, and more. To add or remove options for a customized navigation pane, use the **Customize navigation** control on the navigation pane. This opens the **Customize your navigation pane** settings so you can configure which items appear in the navigation pane.

Expand table

Navigation	Comments
	Select **Home** to return to the compliance portal main page. Visit **Compliance Manager** to check your compliance score and start managing compliance for your organization. Select the **Data classification** section to access trainable classifiers, Sensitive information type entity definitions, content and activity explorers. Select **Data connectors** to configure connectors to import and archive data in your Microsoft 365 subscription. Go to **Alerts** to view and resolve alerts Go to **Policies** to set up policies to govern data, manage devices, and receive alerts. You can also access your DLP and retention policies. Select **Roles & scopes** to manage who in your organization has access to the compliance portal to view content and complete tasks with Permissions and Adaptive scopes. Select **Trials** to discover more security and compliance capabilities in Microsoft 365 Defender and Microsoft Purview for free. Starting a trial helps you make informed decisions about upgrades and purchases. During the free trials, you'll have access to unlimited licenses. You can end trials at any time. Use the links in the **Solutions** section to access your organization's compliance solutions. These include: Catalog Discover, learn about, and start using the intelligent compliance and risk management solutions available to your organization. Audit Use the Audit log to investigate common support and compliance issues. Content search Use Content search to quickly find email in Exchange mailboxes, documents in SharePoint sites and OneDrive locations, and instant messaging conversations in Microsoft Teams and Skype for Business.

Navigation	Comments
	Communication compliance Minimize communication risks by automatically capturing inappropriate messages, investigating possible policy violations, and taking steps to remediate. Data Loss Prevention Detect sensitive content as it's used and shared throughout your organization, in the cloud and on devices, and helps prevent accidental data loss. eDiscovery Expand this section to use the core and eDiscovery (Premium) for preserving, collecting, reviewing, analyzing, and exporting content that's responsive to your organization's internal and external investigations. Data lifecycle management Manage the lifecycle of sensitive data by using features to import, store, and classify business-critical data so you can keep what you need and delete what you don't. Information protection Discover, classify, and protect sensitive and business-critical data throughout its lifecycle across your organization. Information barriers Restrict two-way communication and collaboration between groups and users in Microsoft Teams, SharePoint, and OneDrive. Insider risk management Detect risky activity across your organization to help you quickly identify, investigate, and take action on insider risks and threats. Records management Manage retention and deletion of high-value items for business, legal, or regulatory record-keeping requirements. Privacy risk management Proactively identify and protect against privacy risks such as data hoarding, problematic data transfers, and data oversharing with Microsoft Priva. Subject rights requests Help alleviate the complexity and length of time involved in responding to data subject inquires with Microsoft Priva.

How do I access the compliance portal?

To access the compliance portal, go to https://compliance.microsoft.com and sign in as a global administrator, compliance administrator, or compliance data administrator.

Next steps

- **Visit Microsoft Purview Compliance Manager** to see your compliance score and start managing compliance for your organization. To learn more, see Compliance Manager.
- **Configure insider risk management policies** to help minimize internal risks and enable you to detect, investigate, and take action for risky activities in your organization. See Learn about insider risk management.
- **Review your organization's data loss prevention policies** and make required changes as necessary. To learn more about, see Learn about data loss prevention.
- **Get acquainted with and set up Microsoft Defender for Cloud Apps**. See Quickstart: Get started with Microsoft Defender for Cloud Apps.
- **Learn about and create communication compliance policies** to quickly identify and remediate corporate code-of-conduct policy violations. See Learn about communication compliance.
- **Visit your compliance portal often**, and make sure to review any alerts or potential risks that arise. Go to https://compliance.microsoft.com and sign in.

Feedback

Was this page helpful? Yes No

Identity and access management overview

Article • 05/18/2023

How do Microsoft online services protect production systems from unauthorized or malicious access?

Microsoft online services are designed to allow Microsoft's engineers to operate services without accessing customer content. By default, Microsoft engineers have Zero Standing Access (ZSA) to customer content and no privileged access to the production environment. Microsoft online services use a Just-In-Time (JIT), Just-Enough-Access (JEA) model to provide service team engineers with temporary privileged access to production environments when such access is required to support Microsoft online services. The JIT access model replaces traditional, persistent administrative access with a process for engineers to request temporary elevation into privileged roles when required.

Engineers assigned to a service team to support production services request eligibility for a service team account through an identity and access management solution. The request for eligibility triggers a series of personnel checks to ensure the engineer has passed all cloud screening requirements, completed necessary training, and received appropriate management approval prior to account creation. Only after meeting all eligibility requirements can a service team account be created for the requested environment. To maintain eligibility for a service team account, personnel must go through role-based training annually and rescreening every two years. Failure to complete or pass these checks result in eligibilities automatically being revoked.

Service team accounts don’t grant any standing administrator privileges or access to customer content. When an engineer requires additional access to support Microsoft online services, they request temporary elevated access to the resources they require using an access management tool called Lockbox. Lockbox restricts elevated access to the minimum privileges, resources, and time needed to complete the assigned task. If an authorized reviewer approves the JIT access request, the engineer is granted temporary access with only the privileges necessary to complete their assigned work. This temporary access requires multi-factor authentication and is automatically revoked after the approved period expires.

JEA is enforced by eligibilities and Lockbox roles at the time of request for JIT access. Only requests for access to assets within the scope of the engineer's eligibilities are accepted and passed on to the approver. Lockbox automatically rejects JIT requests that are outside the scope of the engineer's eligibilities and Lockbox roles, including requests that exceed allowed thresholds.

How do Microsoft online services use role-based access control (RBAC) with Lockbox to enforce least privilege?

Service team accounts don't grant any standing administrator privileges or access to customer content. JIT requests for limited administrator privileges are managed through Lockbox. Lockbox uses RBAC to limit the types of JIT elevation requests engineers can make, providing an additional layer of protection to enforce least privilege. RBAC also helps enforce separation of duties by limiting service team accounts to appropriate roles. Engineers supporting a service are granted membership to security groups based on their role. Membership in a security group doesn't grant any privileged access. Instead, security groups allow engineers to use Lockbox to request JIT elevation when required for supporting the system. The specific JIT requests an engineer can make are limited by their security group memberships.

How do Microsoft online services handle remote access to production systems?

Microsoft online services system components are housed in datacenters geographically separated from the operations teams. Datacenter personnel don't have logical access to Microsoft online services systems. As a result, Microsoft service team personnel manage the environment through remote access. Service team personnel who require remote access to support Microsoft online services are only granted remote access after approval from an authorized manager. All remote access uses FIPS 140-2 compatible TLS for secure remote connections.

Microsoft online services use Secure Admin Workstations (SAW) for service team remote access to help protect Microsoft online service environments from compromise. These workstations are designed to prevent intentional or unintentional loss of production data, including locking down USB ports and limiting the software available on the Secure Admin Workstation to what is required for supporting the environment. Secure Admin Workstations are closely tracked and monitored to detect and prevent malicious or inadvertent compromise of customer data by Microsoft engineers.

Privileged access by Microsoft personnel follows a specific path through Microsoft-controlled TSGs with two-factor authentication. All access and activities through TSGs are closely monitored, and alerts and reports are used to identify any anomalous connections. Service teams also implement trend-based monitoring to ensure service health and detect abnormal usage patterns.

How does Customer Lockbox add additional protection for customer content?

Customers can add an additional level of access control to their content by enabling Customer Lockbox. When a Lockbox elevation request involves access to customer content, Customer Lockbox requires approval from the customer as a final step in the approval workflow. This process gives organizations the option to approve or deny these requests and provides direct access control to the customer. If the customer rejects a Customer Lockbox request, access to the requested content is denied. If the customer doesn't reject or approve the request within a certain period, then the request will expire automatically without Microsoft obtaining access to customer content. If the customer approves the request, then Microsoft's temporary access to customer content will be logged, auditable, and revoked automatically after the time assigned to complete the troubleshooting operation expires.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to identity and access control.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
ISO 27001/27002 Statement of Applicability Certificate	A.9.1: Business requirements of access control A.9.2: User access management A.9.3: User responsibilities A.9.4: System and application access control A.15.1: Information security in supplier relationships	April 24, 2023

External audits	Section	Latest report date
ISO 27017 Statement of Applicability Certificate	A.9.1: Business requirements of access control A.9.2: User access management A.9.3: User responsibilities A.9.4: System and application access control A.15.1: Information security in supplier relationships	April 24, 2023
SOC 1 SOC 2 SOC 3	OA-2: Provisioning access OA-7: JIT access OA-21: Secure Admin Workstations and MFA	November 17, 2023

Microsoft 365

Expand table

External audits	Section	Latest report date
FedRAMP (Office 365)	AC-2: Account management AC-3: Access enforcement AC-5: Separation of duties AC-6: Least privilege AC-17: Remote access	July 31, 2023
ISO 27001/27002/27017 Statement of Applicability Certification (27001/27002) Certification (27017)	A.9.1: Business requirements of access control A.9.2: User access management A.9.3: User responsibilities A.9.4: System and application access control A.15.1: Information security in supplier relationships	March 2023
SOC 1	CA-33: Account modification CA-34: User authentication CA-35: Privileged access CA-36: Remote access CA-57: Customer Lockbox Microsoft management approval CA-58: Customer Lockbox service requests CA-59: Customer Lockbox notifications CA-61: JIT review and approval	January 3, 2023
SOC 2	CA-32: Shared account policy CA-33: Account modification CA-34: User authentication CA-35: Privileged access	January 3, 2023

External audits	Section	Latest report date
	CA-36: Remote access CA-53: Third-party monitoring CA-56: Customer Lockbox customer approval CA-57: Customer Lockbox Microsoft management approval CA-58: Customer Lockbox service requests CA-59: Customer Lockbox notifications CA-61: JIT review and approval	
SOC 3	CUEC-15: Customer Lockbox requests	January 3, 2023

Resources

- Video: Microsoft 365 access controls

Feedback

Was this page helpful? Yes No

Customer Lockbox for Microsoft Azure

Article • 10/17/2023

> **Note**
>
> To use this feature, your organization must have an **Azure support plan** with a minimal level of **Developer**.

Most operations, support, and troubleshooting performed by Microsoft personnel and sub-processors do not require access to customer data. In those rare circumstances where such access is required, Customer Lockbox for Microsoft Azure provides an interface for customers to review and approve or reject customer data access requests. It is used in cases where a Microsoft engineer needs to access customer data, whether in response to a customer-initiated support ticket or a problem identified by Microsoft.

This article covers how to enable Customer Lockbox and how Lockbox requests are initiated, tracked, and stored for later reviews and audits.

Supported services

The following services are currently supported for Customer Lockbox:

- Azure API Management
- Azure App Service
- Azure Cognitive Search
- Azure Cognitive Services
- Azure Container Registry
- Azure Data Box
- Azure Data Explorer
- Azure Data Factory
- Azure Data Manager for Energy
- Azure Database for MySQL
- Azure Database for MySQL Flexible Server
- Azure Database for PostgreSQL
- Azure Databricks
- Azure Edge Zone Platform Storage
- Azure Energy
- Azure Functions
- Azure HDInsight

- Azure Health Bot
- Azure Intelligent Recommendations
- Azure Kubernetes Service
- Azure Load Testing (CloudNative Testing)
- Azure Logic Apps
- Azure Monitor
- Azure Red Hat OpenShift
- Azure Spring Apps
- Azure SQL Database
- Azure SQL Managed Instance
- Azure Storage
- Azure Subscription Transfers
- Azure Synapse Analytics
- Azure Unified Vision Service
- Commerce AI (Intelligent Recommendations)
- DevCenter / DevBox
- ElasticSan
- Kusto (Dashboards)
- Microsoft Azure Attestation
- OpenAI
- Spring Cloud
- Unified Vision Service
- Virtual Machines in Azure

Enable Customer Lockbox

You can now enable Customer Lockbox from the Administration module in the Customer Lockbox blade.

> **Note**
>
> To enable Customer Lockbox, the user account needs to have the **Global Administrator role assigned**.

Workflow

The following steps outline a typical workflow for a Customer Lockbox request.

1. Someone at an organization has an issue with their Azure workload.

2. After this person troubleshoots the issue, but can't fix it, they open a support ticket from the Azure portal. The ticket is assigned to an Azure Customer Support Engineer.

3. An Azure Support Engineer reviews the service request and determines the next steps to resolve the issue.

4. If the support engineer can't troubleshoot the issue by using standard tools and service generated data, the next step is to request elevated permissions by using a Just-In-Time (JIT) access service. This request can be from the original support engineer or from a different engineer because the problem is escalated to the Azure DevOps team.

5. After the access request is submitted by the Azure Engineer, Just-In-Time service evaluates the request taking into account factors such as:

 - The scope of the resource
 - Whether the requester is an isolated identity or using multifactor authentication
 - Permissions levels Based on the JIT rule, this request might also include an approval from Internal Microsoft Approvers. For example, the approver might be the Customer support lead or the DevOps Manager.

6. When the request requires direct access to customer data, a Customer Lockbox request is initiated. For example, remote desktop access to a customer's virtual machine.

 The request is now in a **Customer Notified** state, waiting for the customer's approval before granting access.

7. The approver(s) at the customer organization for a given Lockbox request are determined as follows:

 - For Subscription scoped requests (requests to access specific resources contained within a subscription), users who have been assigned the Owner role on the associated subscription.
 - For Tenant scope requests (requests to access the Microsoft Entra tenant), users who have been assigned the Global Administrator role on the Tenant.

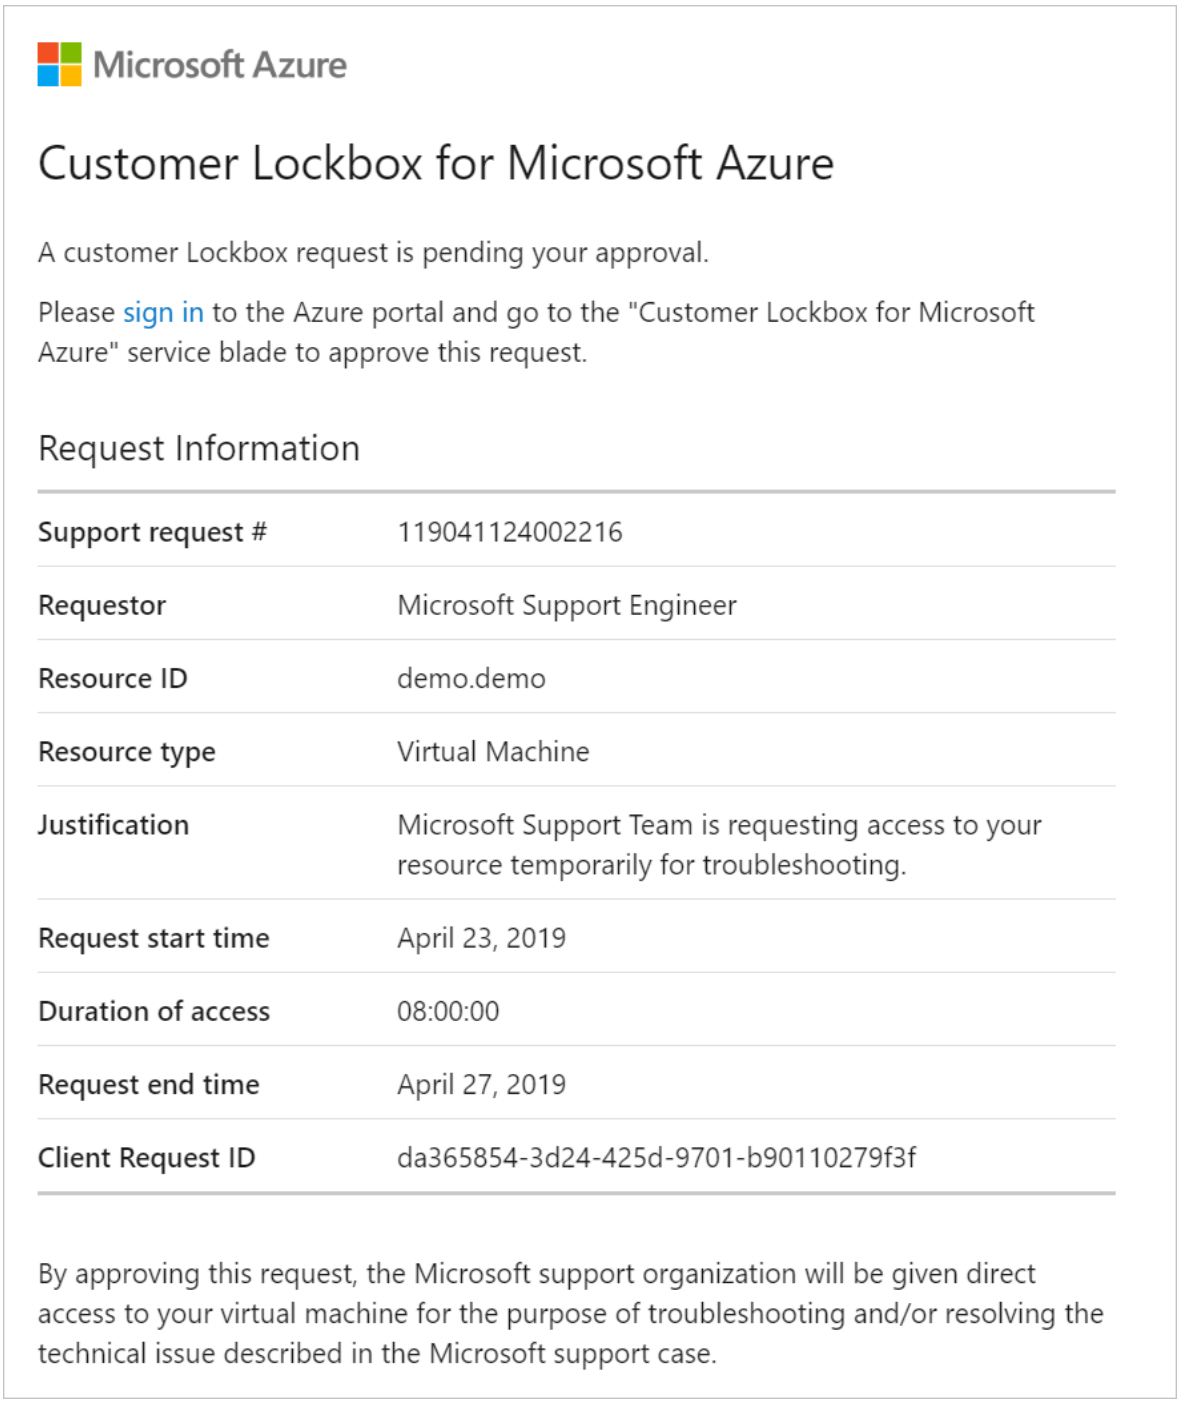
 > **Note**
 >
 > Role assignments must be in place before Lockbox starts to process a request. Any role assignments made after Lockbox starts to process a given request will not be recognized by Lockbox. Because of this, to use PIM eligible

assignments for the Subscription Owner role, users are required to activate the role before the Customer Lockbox request is initiated. Refer to **Activate Microsoft Entra roles in PIM / Activate Azure resource roles in PIM** for more information on activating PIM eligible roles.

Role assignments scoped to management groups are not supported in Lockbox at this time.

8. At the customer organization, designated lockbox approvers (Azure Subscription Owner/Microsoft Entra Global admin receive an email from Microsoft to notify them about the pending access request.

 Example email:



Microsoft Azure

Customer Lockbox for Microsoft Azure

A customer Lockbox request is pending your approval.

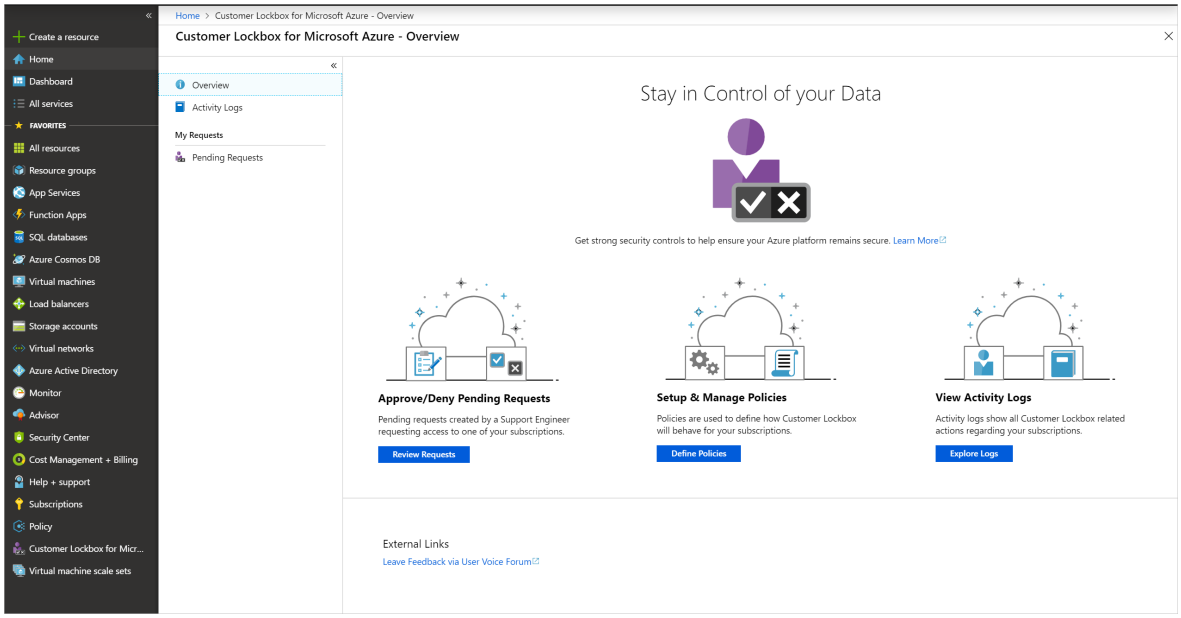
Please sign in to the Azure portal and go to the "Customer Lockbox for Microsoft Azure" service blade to approve this request.

Request Information

Support request #	119041124002216
Requestor	Microsoft Support Engineer
Resource ID	demo.demo
Resource type	Virtual Machine
Justification	Microsoft Support Team is requesting access to your resource temporarily for troubleshooting.
Request start time	April 23, 2019
Duration of access	08:00:00
Request end time	April 27, 2019
Client Request ID	da365854-3d24-425d-9701-b90110279f3f

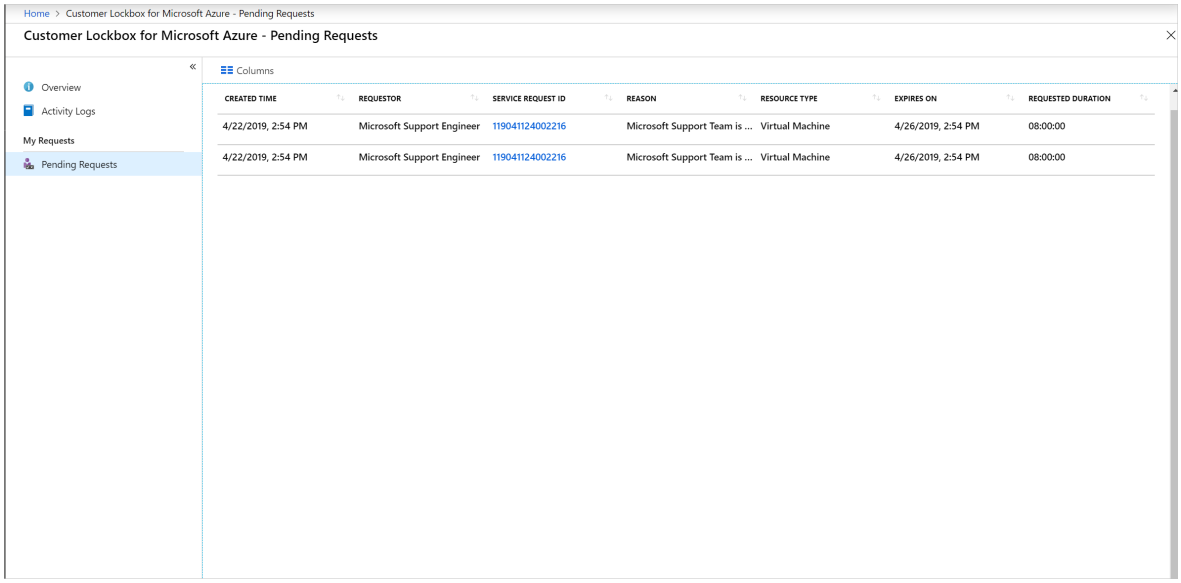
By approving this request, the Microsoft support organization will be given direct access to your virtual machine for the purpose of troubleshooting and/or resolving the technical issue described in the Microsoft support case.

9. The email notification provides a link to the **Customer Lockbox** blade in the Administration module. Using this link, the designated approver signs in to the Azure portal to view any pending requests that their organization has for Customer Lockbox:



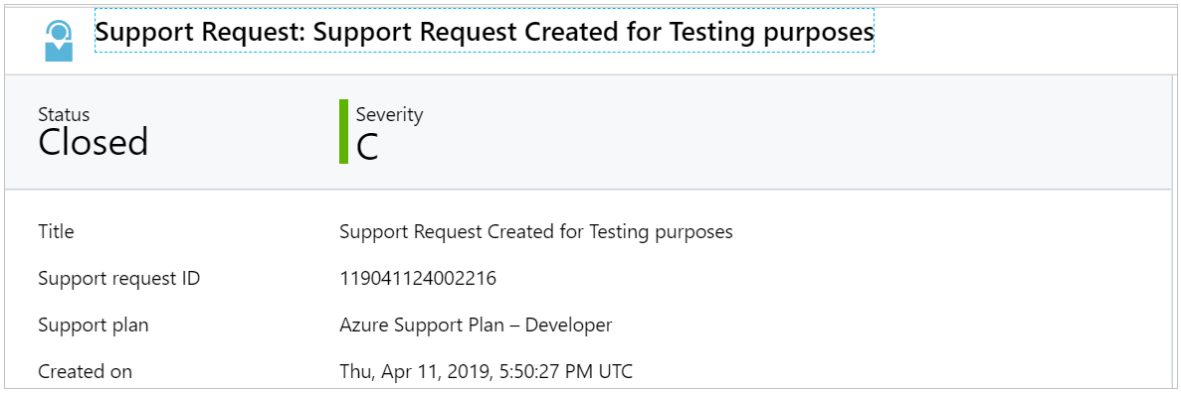
The request remains in the customer queue for four days. After this time, the access request automatically expires and no access is granted to Microsoft engineers.

10. To get the details of the pending request, the designated approver can select the lockbox request from **Pending Requests**:

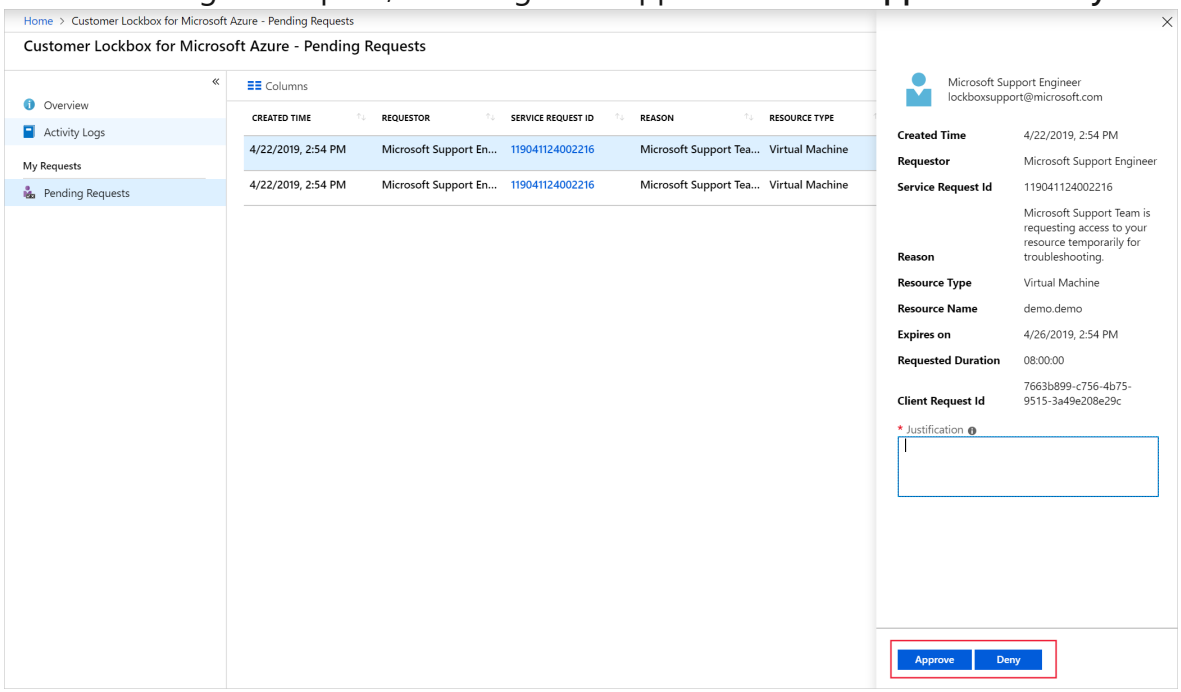


11. The designated approver can also select the **SERVICE REQUEST ID** to view the support ticket request that was created by the original user. This information provides context for why Microsoft Support is engaged, and the history of the

reported problem. For example:



12. After reviewing the request, the designated approver selects **Approve** or **Deny**:



As a result of the selection:

- **Approve**: Access is granted to the Microsoft engineer. The access is granted for a default period of eight hours.
- **Deny**: The elevated access request by the Microsoft engineer is rejected and no further action is taken.

For auditing purposes, the actions taken in this workflow are logged in Customer Lockbox request logs.

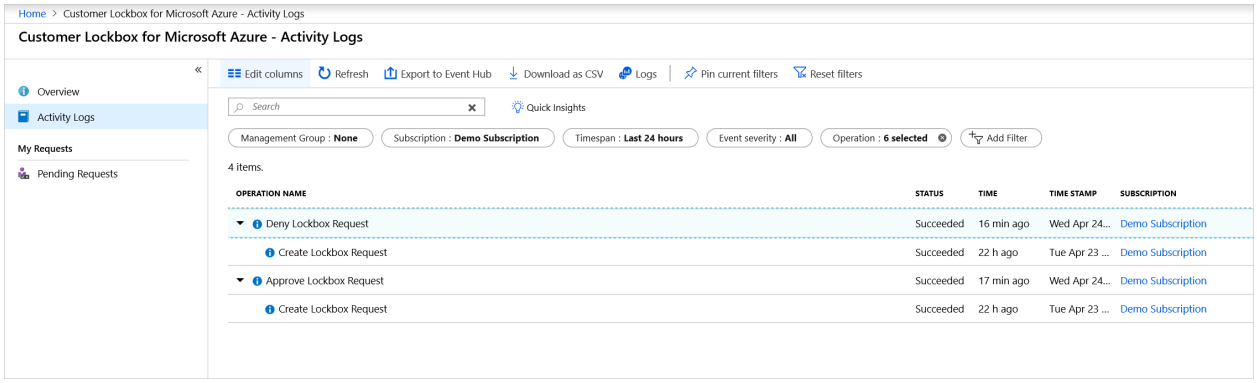
Auditing logs

Customer Lockbox logs are stored in activity logs. In the Azure portal, select **Activity Logs** to view auditing information related to Customer Lockbox requests. You can filter for specific actions, such as:

- **Deny Lockbox Request**

- **Create Lockbox Request**
- **Approve Lockbox Request**
- **Lockbox Request Expiry**

As an example:



Customer Lockbox integration with the Microsoft cloud security benchmark

We've introduced a new baseline control (PA-8: Determine access process for cloud provider support) in the Microsoft cloud security benchmark that covers Customer Lockbox applicability. Customers can now leverage the benchmark to review Customer Lockbox applicability for a service.

Exclusions

Customer Lockbox requests are not triggered in the following engineering support scenarios:

- Emergency scenarios that fall outside of standard operating procedures. For example, a major service outage requires immediate attention to recover or restore services in an unexpected or unpredictable scenario. These “break glass” events are rare and, in most instances, do not require any access to customer data to resolve.
- A Microsoft engineer accesses the Azure platform as part of troubleshooting and is inadvertently exposed to customer data. For example, the Azure Network Team performs troubleshooting that results in a packet capture on a network device. It is rare that such scenarios would result in access to meaningful quantities of customer data. Customers can further protect their data through use of in transit and at rest encryption.

Customer Lockbox requests are also not triggered by external legal demands for data. For details, see the discussion of government requests for data on the Microsoft Trust

Center.

Next steps

Customer Lockbox is available for all customers who have an Azure support plan with a minimal level of **Developer**. You can enable Customer Lockbox from the Administration module in the Customer Lockbox blade.

Customer Lockbox requests are initiated by a Microsoft engineer if this action is needed to progress a support case.

Account management in Microsoft 365

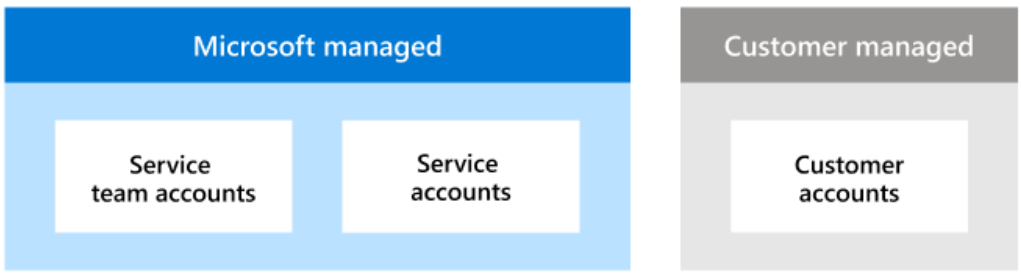
Article • 03/02/2023

Microsoft has invested heavily in systems and controls that automate most Microsoft 365 operations while intentionally limiting the need for direct access to servers and customer data by service personnel. Humans govern the service and software operates the service. This structure enables Microsoft to manage Microsoft 365 at scale and minimizes the risks of both internal and external threats. Microsoft approaches access control with the assumption that everyone is a potential threat to Microsoft 365 services and customer data. For this reason, the Zero Standing Access (ZSA) principle lays the foundation for the entire access control structure used by Microsoft 365.

By default, Microsoft personnel have zero standing privileged access to any Microsoft 365 environment or customer data for an organization. It is only through a robust system of checks and approvals that service team personnel can gain privileged access with a narrow action and time scope. Through this system, Microsoft can significantly reduce the potential of Microsoft 365 service personnel and attackers from gaining unauthorized access or causing malicious or accidental harm to Microsoft services and customers.

Account types

Microsoft 365 meets all organizational missions and business functions using three categories of accounts: service team accounts, service accounts, and customer accounts. Managing these accounts is a shared responsibility between Microsoft and customers. Microsoft manages both service team and service accounts, which are used to operate and support Microsoft products and services. Customer accounts are managed by the customer and allow them to tailor account access to meet their internal access control requirements.



Microsoft-managed accounts

Service team accounts are used by Microsoft 365 service team personnel developing and maintaining Microsoft 365 services. These accounts do not have standing privileged access to Microsoft 365 services, instead they can be used to request temporary and limited privileged access to perform a specified job function. Not every service team account can perform the same actions, separation of duties is enforced using role-based access control (RBAC). Roles ensure that service team members and their accounts have only the minimum access required to perform specific job duties. Additionally, service team accounts cannot belong to multiple roles where they can act as the approver for their own actions.

Service accounts are used by Microsoft 365 services to authenticate when communicating with other services through automated processes. Just as service team accounts are only given the minimum access necessary to perform the specific personnel's job duties, service accounts are only granted the bare minimum access needed for their intended purpose. Additionally, there are multiple types of service accounts that are designed to fulfill a specific need. One Microsoft 365 service may have multiple service accounts, each with a different role to perform.

Customer-managed accounts

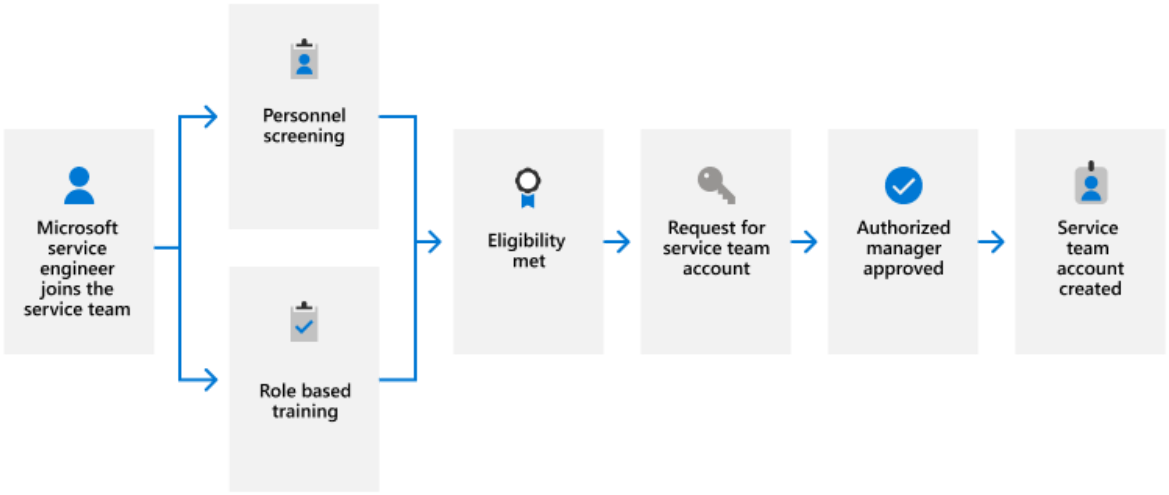
Customer accounts are used to access Microsoft 365 service and are the only accounts each customer is responsible for. It is the customer's duty to create and manage the accounts in their organization to maintain a secure environment. Management of customer accounts is done through Microsoft Entra ID or federated with on-premises Active Directory (AD). Each customer has a unique set of access control requirements they must meet, and customer accounts grant each customer the ability to satisfy their individual needs. Customer accounts cannot access any data outside of their customer organization.

Service team account management

Microsoft 365 manages service team accounts throughout their lifecycle using an account management system called Identity Management (IDM). IDM uses a combination of automated verification processes and managerial approval to enforce the security requirements related to service team account access.

Service team members do not automatically get a service team account, they must first meet eligibility requirements and get approval from an authorized manager. To be

eligible for a service team account, service team personnel must first go through pre-employment personnel screening, a cloud background check, and complete all standard and required role-based training. Once all eligibility requirements have been met, a request for a service team account can be made and must be approved by an authorized manager.



IDM is also responsible for tracking the periodic rescreening and training needed to maintain a service team account. The Microsoft cloud background check must be completed every two years and all training material must be reviewed annually. If either of these requirements are not satisfied by the expiration date, their eligibility is revoked and the service team account is automatically disabled.

Additionally, service team account eligibility is automatically updated by personnel transfer and termination. Changes made in the Human Resources Information System (HRIS) trigger IDM to take action, which varies depending on the situation. Personnel transferring to another service team will have an expiration date set for their eligibilities and a request to maintain eligibilities must be submitted by the service team member and approved by their new manager. Terminated personnel automatically have all eligibilities revoked and their service team account is disabled on their last day. An urgent request for account revocation can be made for involuntary terminations.

By default, Service team accounts have limited read access to broad system metadata used for regular troubleshooting. Additionally, baseline service team accounts cannot request privileged access to Microsoft 365 or customer data. Another request must be made for the service team account to be added to a role that allows for the service team member to request elevated privileges to perform specific tasks and operations.

Microsoft 365 service engineer access control

Article • 05/25/2023

Zero Standing Access (ZSA) means that Microsoft service team personnel don't have any standing privileged access to the Microsoft 365 production environment or to customer data. When a Microsoft service team member wants to update the service or access customer data for any reason, they must submit a request justifying the need and receive approval from an authorized manager. At scale, it isn't feasible to manually provide and remove access as required to maintain Microsoft 365 services, so Microsoft has developed an automated solution to manage privileged access as needed.

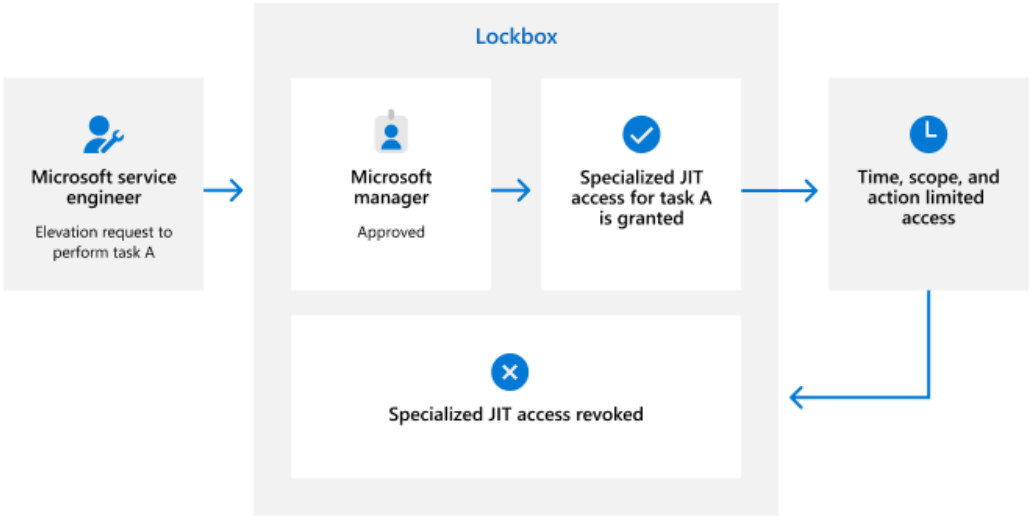
Lockbox

All access to Microsoft 365 systems and customer data is brokered by Lockbox, an access control system that uses a Just-In-Time (JIT) and Just-Enough-Access (JEA) model to provide service engineers with temporary privileged access to specified Microsoft 365 services and data. Additionally, all requests and actions are logged for auditing purposes and are accessible by using the Office 365 Management Activity API and the Security and Compliance Center.

Before a Microsoft service engineer can connect to any Microsoft 365 systems or access customer data, they must submit an access request through Lockbox. This request can only be approved if certain criteria are met:

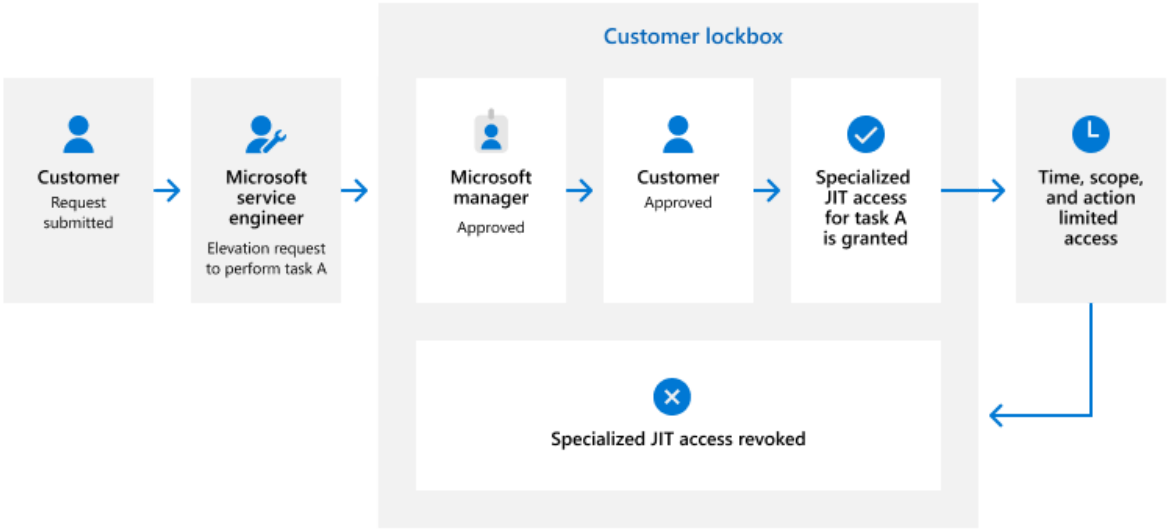
- The service engineer meets the eligibility requirements for a service team account,
- They belong to the Lockbox role associated with the work in the request,
- The requested access time doesn't exceed the maximum permitted time,
- They have a legitimate business justification,
- The requested resource they wish to access is within their work scope, and
- They receive manager approval

Once all criteria have been met and verified by Lockbox, temporary access is granted to perform the specific action requested. After the time for the request has elapsed, the access is revoked.



Additionally, if a customer licenses and enables the Customer Lockbox feature, any attempt from a Microsoft service engineer to access customer data must be additionally approved by an administrator in the customer tenant. The need to access customer data can arise from both customer and Microsoft. For example, an incident raised by a customer may require access to their data to fix the issue or when Microsoft needs data access to apply a specific update.

Customers don't have any tools to initiate a Customer Lockbox request; they must submit a ticket to Microsoft that requires a Customer Lockbox request to be raised. A Customer Lockbox request raised by a Microsoft service engineer must be approved by a Microsoft Manager and an authorized administrator in the customer tenant.



Lockbox roles

To enforce separation of duties and the principle of least privilege, service engineers must belong to a Lockbox role that corresponds to their role on the team. Lockbox roles are managed within the Identity Management tool and define the privileges and actions a service team member can be approved for through the Lockbox request process. Service team personnel must request to be a member of a Lockbox role and receive managerial approval. If approved, the employee's service team account is placed in a security group enforced by Active Directory (AD) and Microsoft Entra ID.

Constrained management interfaces

Service engineers use two management interfaces to perform administrative tasks: Remote Desktop from a Secure Admin Workstation (SAW) through a secured Terminal Service Gateway (TSG) and Remote PowerShell. Within these management interfaces, access controls based on the approved Lockbox request and software policies place significant restrictions on which applications are executed and which commands and cmdlets are available.

Remote Desktop

Service team members administrating their service using Remote Desktop must connect from a SAW, specially designed and manufactured laptops managed by Microsoft specifically for this use case. Microsoft partners with suppliers to build SAWs, creating a short and secure supply chain. SAWs use hardened operating systems that are configured to limit all functionality except what is needed for defined management tasks. These limitations include disabling of all USB ports, strict application access lists, removal of email access, limiting internet browsing, and enforcing inactivity screensaver lockouts. Microsoft access control systems examine SAW machines periodically to ensure they're compliant with the latest security controls and automatically disables machines if they're determined to be non-compliant.

Service engineers are only permitted to connect to one TSG at a time and multiple sessions aren't allowed. However, TSGs allow Microsoft 365 service team administrators to connect to multiple servers, each with only one concurrent session, so that administrators can effectively perform their duties. Service team administrators don't have any permissions on the TSGs themselves. The TSG is used only to enforce multifactor authentication (MFA) and encryption requirements. Once the service team administrator connects to a specific server through a TSG, the specific server enforces a session limit of one per administrator.

Usage restrictions, connection, and configuration requirements for Microsoft 365 personnel are established by Active Directory group policies. These policies include the

following TSG characteristics:

- Use only FIPS 140-2 validated encryption
- Sessions disconnect after 15 minutes of inactivity
- Sessions automatically sign out after 24 hours

Connections to TSGs also require MFA using a separate physical smart card. Service engineers are issued different smart cards for various platforms and secrets management platforms ensure secure storage of credentials. TSGs use Active Directory group policies to control who can log into remote servers, the number of allowed sessions, and idle timeout settings.

Remote PowerShell

In addition to remote access using specially configured TSGs, service team personnel with the Service Engineer Operations Lockbox role can access certain administrative functionality on production servers using Remote PowerShell. To use this access, the user must be authorized for read-only (debug) access to the Microsoft 365 production environment. Privilege escalation is enabled the same way it's enabled for TSGs using the Lockbox process.

For remote access, each datacenter has a load-balanced virtual IP that serves as a single point of access. The available Remote PowerShell cmdlets are based on the privilege level identified in the access claim obtained during authentication. These cmdlets provide the only administrative functionality accessible by users connecting using this method. Remote PowerShell limits the scope of commands available to the engineer and is based on the level of access granted via the Lockbox process. For example, in Exchange Online, the Get-Mailbox cmdlet might be available, but the Set-Mailbox cmdlet wouldn't.

Resources

- Isolation in Microsoft 365
- Microsoft pre-employment screening
- Microsoft cloud background check

Microsoft Purview Customer Lockbox

Article • 11/15/2023

This article provides deployment and configuration guidance for Customer Lockbox. Customer Lockbox supports requests to access data in Exchange Online, SharePoint Online, OneDrive for Business, Teams, and Microsoft Copilot for Microsoft 365. To recommend support for other services, submit a request at Feedback Portal.

To see the options for licensing your users to benefit from Microsoft Purview offerings, see the Microsoft 365 licensing guidance for security & compliance.

Customer Lockbox ensures that Microsoft can't access your content to do service operations without your explicit approval. Customer Lockbox brings you into the approval workflow process that Microsoft uses to ensure only authorized requests allow access to your content. To learn more about Microsoft's workflow process, see Privileged access management.

Occasionally, Microsoft engineers help troubleshoot and fix issues that arise with the service. Usually, engineers fix issues using extensive telemetry and debugging tools Microsoft has in place for its services. However, some cases require a Microsoft engineer to access your content to determine the root cause and fix the issue. Customer Lockbox requires the engineer to request access from you as a final step in the approval workflow. This gives you the option to approve or deny the request for your organization, and provide direct-access control to your content.

> **Tip**
>
> If you're not an E5 customer, use the 90-day Microsoft Purview solutions trial to explore how additional Purview capabilities can help your organization manage data security and compliance needs. Start now at the **Microsoft Purview compliance portal trials hub**. Learn details about **signing up and trial terms**.

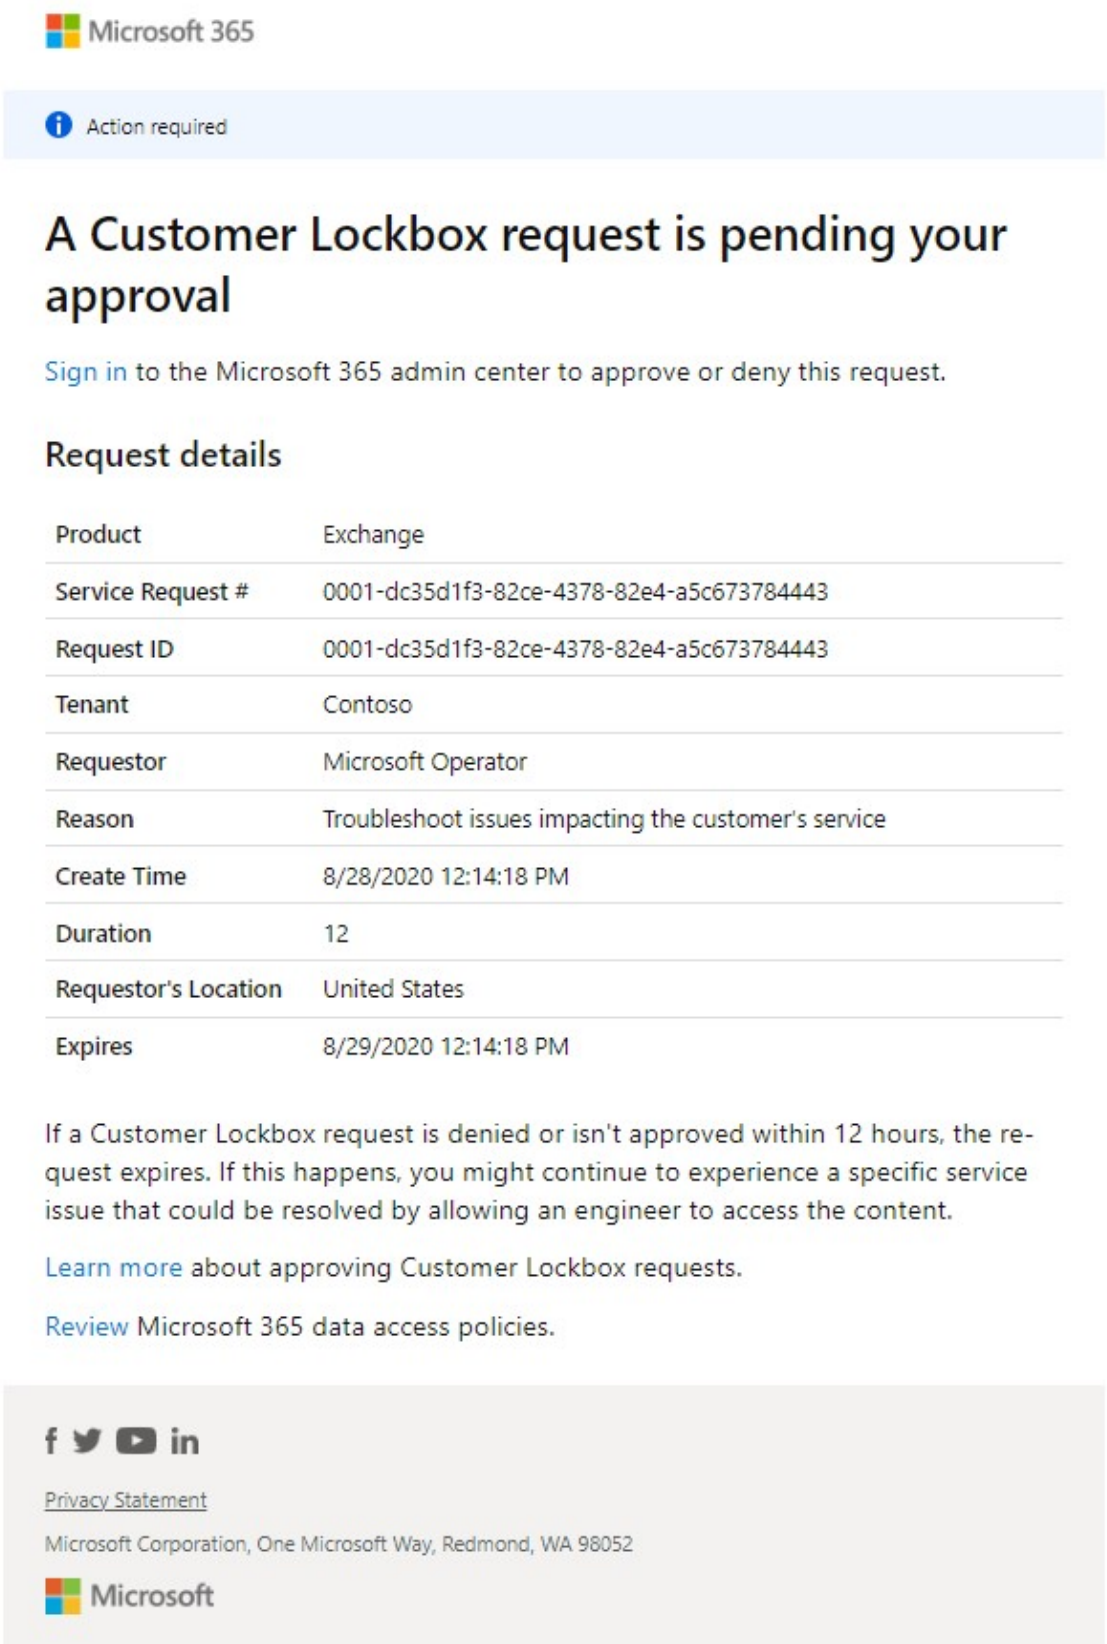
Customer Lockbox overview video

https://www.microsoft.com/en-us/videoplayer/embed/8fecf10b-1f03-4849-8b67-76d3d2a43f26?autoplay=false&postJsllMsg=true

Customer Lockbox workflow

These steps outline the typical workflow when a Microsoft engineer starts a Customer Lockbox request:

1. Someone at an organization experiences an issue with their Microsoft 365 mailbox.
2. After the user troubleshoots the issue, but can't fix it, they open a support request with Microsoft Support.
3. A Microsoft support engineer reviews the service request and determines a need to access the organization's tenant to repair the issue.
4. The Microsoft support engineer logs into the Customer Lockbox request tool and makes a data access request that includes the organization's tenant name, service request number, expected start time of access (starts immediately post-approval if not specified), the estimated amount of time the engineer needs access to the data, and the service the request is for.
5. After a Microsoft Support manager approves the request, Customer Lockbox sends the designated approver at the organization an email notification about the pending access request from Microsoft.



Anyone who is assigned the Customer Lockbox access approver admin role in Microsoft 365 admin center can approve Customer Lockbox requests.

6. The approver signs in to the Microsoft 365 admin center and approves the request. This step triggers the creation of an audit record available by searching the audit log. For more information, see Auditing Customer Lockbox requests.

If the customer rejects the request or doesn't approve the request within 12 hours, the request expires and no access is granted to the Microsoft engineer.

> **Important**
>
> Microsoft does not include any links in Customer Lockbox email notifications requiring you to sign in to Office 365.

7. After the approver from the organization approves the request, the Microsoft engineer receives the approval message, logs into the tenant, and fixes the customer's issue. Microsoft engineers have the requested duration to fix the issue after which the access is automatically revoked.

> **Note**
>
> All actions performed by a Microsoft engineer are logged in the audit log. You can search for and review these audit records.

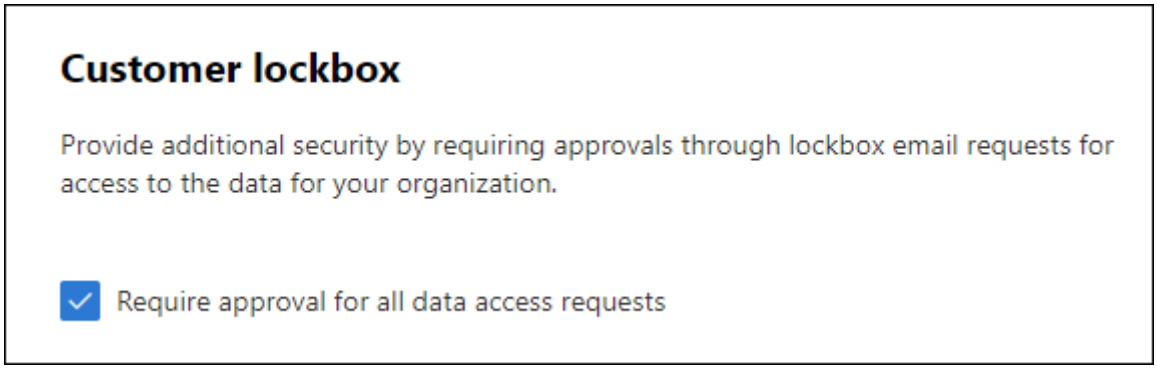
Turn Customer Lockbox requests on or off

You can turn on Customer Lockbox controls in the Microsoft 365 admin center. When you turn on Customer Lockbox, Microsoft must obtain your organization's approval before accessing any of your tenant's content.

1. Using a work or school account that has either the global administrator or the **Customer Lockbox access approver** role assigned, go to https://admin.microsoft.com and sign in.

2. Choose **Settings** > **Org Settings** > **Security & Privacy**.

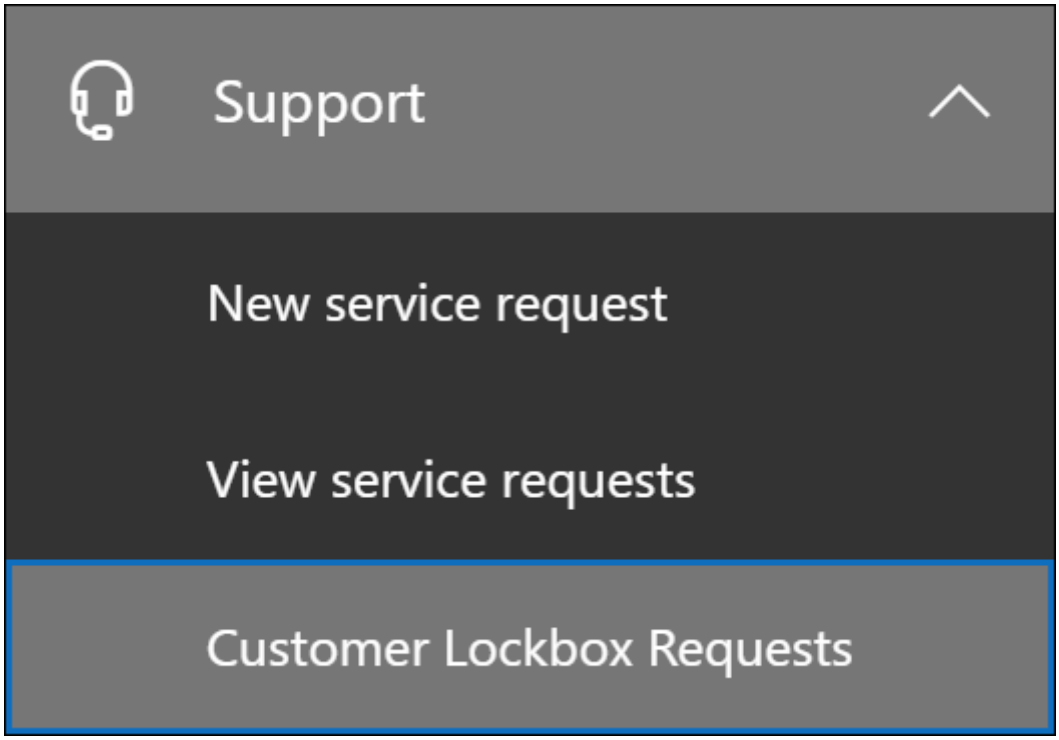
3. Select **Security & Privacy**, then select **Customer Lockbox** in the left column. Check the **Require approval for all data access requests** checkbox and save the changes to turn on the feature.



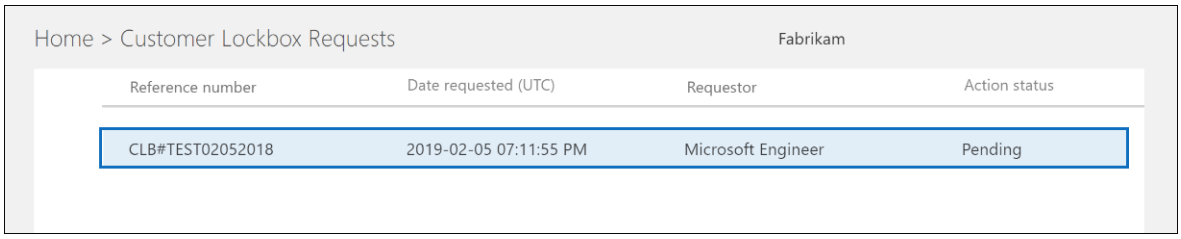
Approve or deny a Customer Lockbox request

1. Using a work or school account that has either the global administrator or the **Customer Lockbox access approver** role assigned, go to https://admin.microsoft.com and sign in.

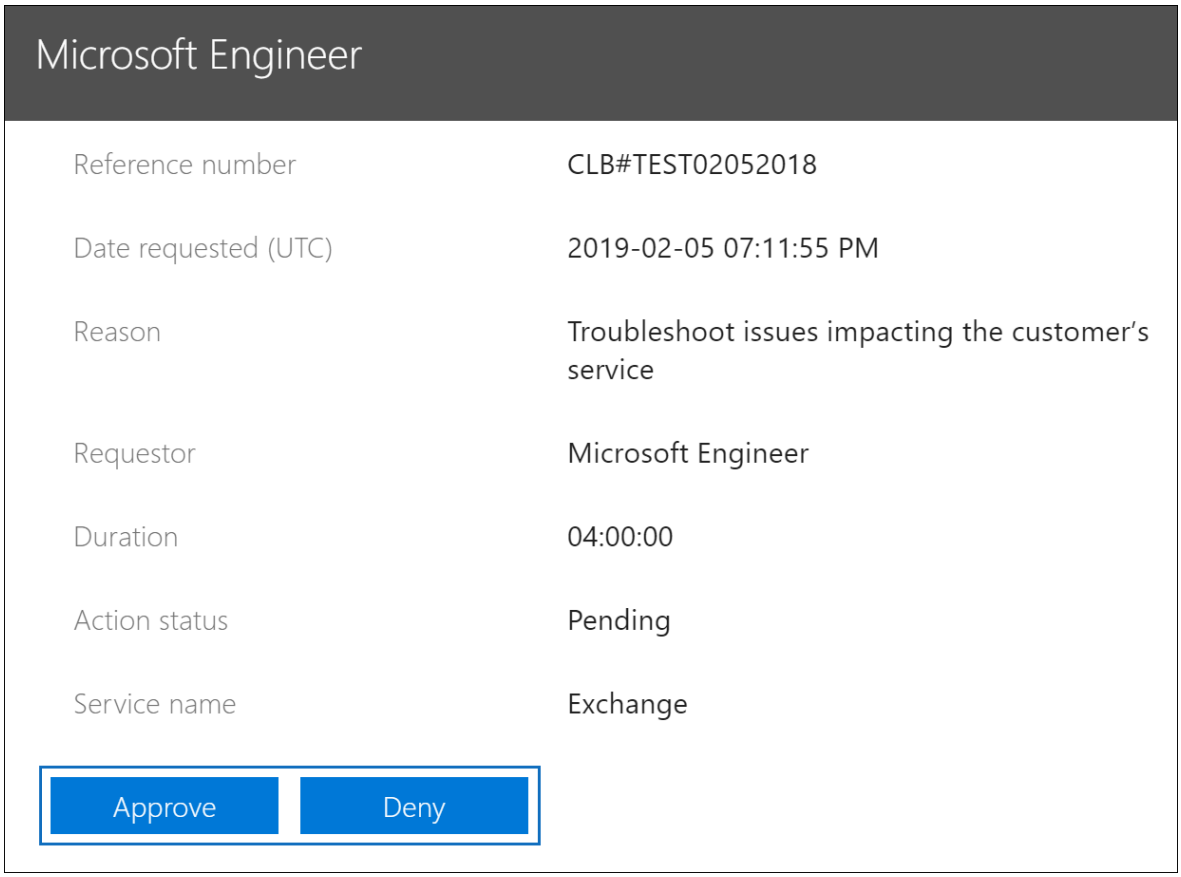
2. Choose **Support** > **Customer Lockbox Requests**.



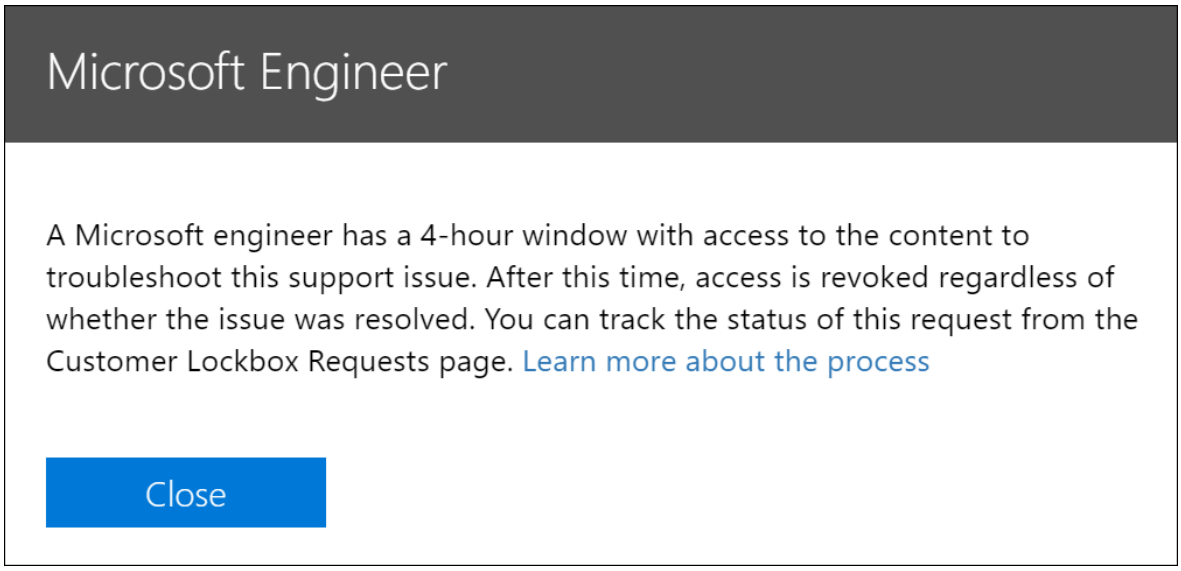
 A list of Customer Lockbox requests displays.



3. Select a Customer Lockbox request, and then choose **Approve** or **Deny**.



A confirmation message about the approval of the Customer Lockbox request displays.



Note

Use the Set-AccessToCustomerDataRequest cmdlet to approve, deny, or cancel Microsoft Purview Customer Lockbox requests that control access to your data by Microsoft support engineers. For more information, see **Set-AccessToCustomerDataRequest**.

Auditing Customer Lockbox requests

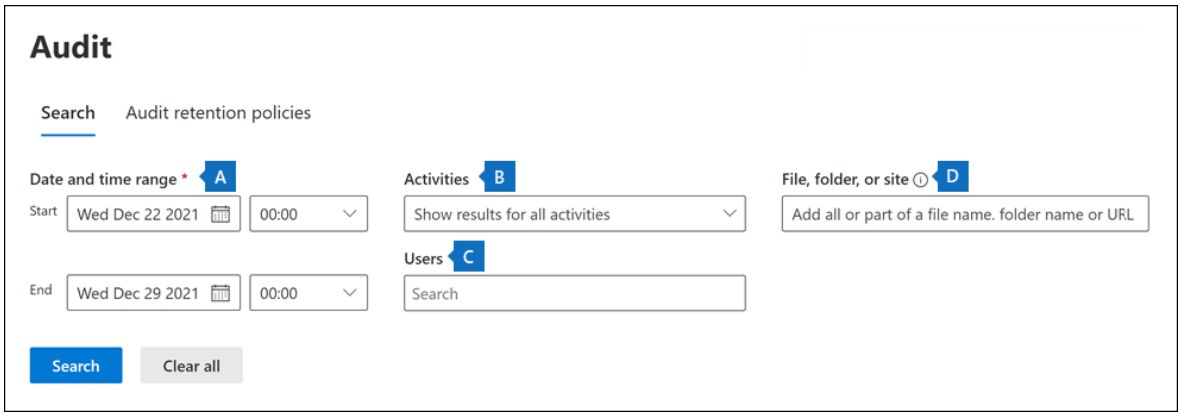
Audit records that correspond to Customer Lockbox requests are logged in the Microsoft 365 audit log. You can access these logs by using the audit log search tool in the Microsoft Purview compliance portal. Actions related to accepting or denying a Customer Lockbox request and actions performed by Microsoft engineers (when access requests are approved) are also logged in the audit log. You can search for and review these audit records.

Search the audit log for activity related to Customer Lockbox requests

Before you can use the audit log to track requests for Customer Lockbox, there are some steps you need to take to set up audit logging, including assigning permissions to search the audit log. For more information, see Set up Microsoft Purview Audit (Standard). Once you've completed the setup, use these steps to create an audit log search query to return audit records related to Customer Lockbox:

1. Go to https://compliance.microsoft.com .
2. Sign in using an account that has been assigned the appropriate permissions to search the audit log.
3. In the left pane of the compliance center, choose **Audit**.

 The **Search** tab on the **Audit** page is displayed.



4. Configure the following search criteria:

 a. **Start date** and **End date**. Select a date and time range to display the events that occurred within that period.

 b. **Activities**. Leave this field blank so that the search returns audit records for all activities. This is necessary to return any audit records related to Customer

Lockbox requests and corresponding activity performed by Microsoft engineers.

c. **Users**. Leave this field blank.

d. **File, folder, or site**. Leave this field blank.

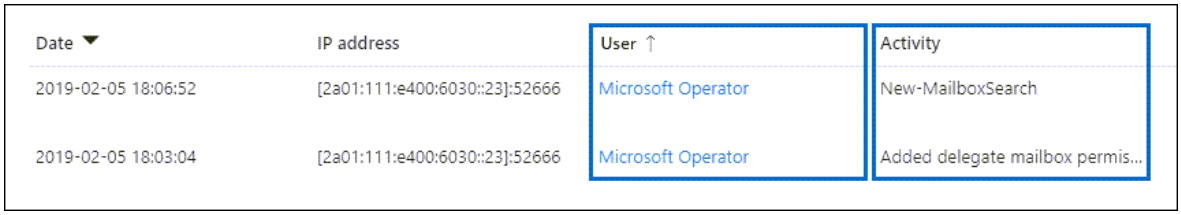
5. Click **Search** to run the search using your search criteria.

 The search results are displayed after a few moments. More search results will be added to the page until the search is complete.

6. Click the header in the **Activity** column to sort the results alphabetically based on the values in the **Activity** column.

7. Scroll down and look for audit records with an activity of **Set-AccessToCustomerDataRequest**. Records with this activity are related to an approver in your organization approving or denying a Customer Lockbox request.

8. Alternatively, click the header in the **User** column to sort the results alphabetically using the values in the **User** column. Look for the value of **Microsoft Operator**, which indicates activities performed by a Microsoft engineer in response to an approved Customer Lockbox request. The **Activity** column displays the action performed by the engineer.



9. In the list of results, click an audit record to display it.

Export the audit log search results

You can also export the audit log search results to a CSV file and then open the file in Excel to use the filtering and sorting capabilities to make it easier to find and view audit records related to a Customer Lockbox access request.

To export audit records, use the previous steps to search the audit log. When the search is complete, select **Export** > **Download all results** at the top of the search results page. When the export process is complete, you can download the CSV file to your local computer. For more detailed instructions, see Export, configure, and view audit log records.

After you download the file, you can open it in Excel and then filter on the **Operations** column to display audit records for **Set-AccessToCustomerDataRequest** activities. You can also filter on the **UserIds** column (using the value **Microsoft Operator**) to display audit records for activities performed by Microsoft engineers.

> **Note**
>
> When viewing audit records in the CSV file, additional information is contained in the **AuditData** column. The information in this column is contained in a JSON object, which contains multiple properties that are configured as *property:value* pairs separated by commas. You can use the JSON transform feature in the Power Query Editor in Excel to split each property in the JSON object in the **AuditData** column into multiple columns so that each property has its own column. This makes it easier to interpret this information. For detail instructions, see **Format the exported audit log using the Power Query Editor**.

Use PowerShell to search and export audit records

An alternative to using the audit search tool in the Microsoft Purview compliance portal is to run the Search-UnifiedAuditLog cmdlet in Exchange Online PowerShell. One advantage of using PowerShell is that you can specifically search for **Set-AccessToCustomerDataRequest** activities or activities performed by Microsoft engineers related to a Customer Lockbox request.

After you connect to Exchange Online PowerShell, run one of the following commands. Replace the placeholders with a specific date range.

Search for `Set-AccessToCustomerDataRequest` activities

PowerShell

```
Search-UnifiedAuditLog -StartDate xx/xx/xxxx -EndDate xx/xx/xxxx -Operations
Set-AccessToCustomerDataRequest
```

Search for activities performed by Microsoft engineers

PowerShell

```
Search-UnifiedAuditLog -StartDate xx/xx/xxxx -EndDate xx/xx/xxxx -UserIds
"Microsoft Operator"
```

For more information and examples, see Use PowerShell to search and export audit log records.

We've also provided a PowerShell script that you can use to search the audit log and export the results to a CSV file. For more information, see Use a PowerShell script to search the audit log.

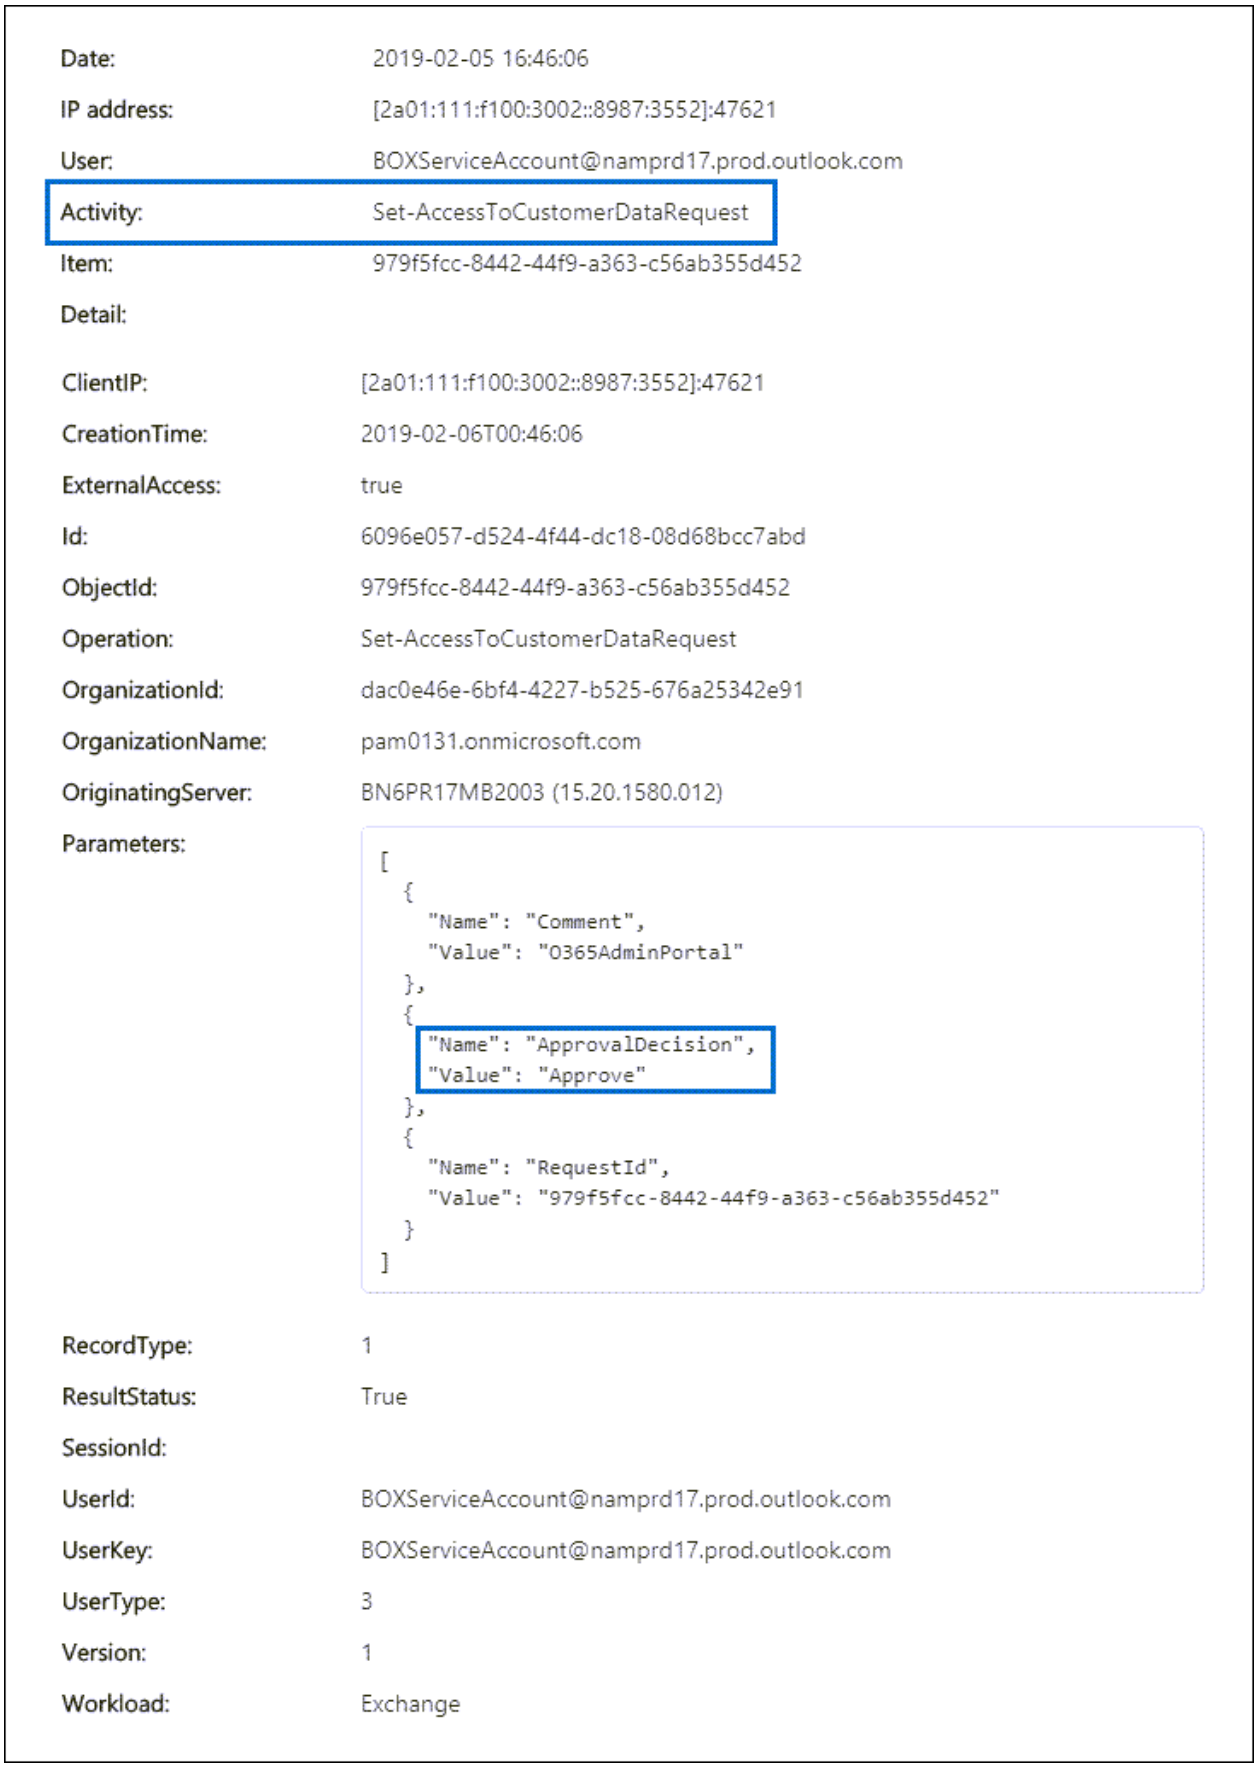
Audit record for a Customer Lockbox request

When a person in your organization approves or denies a Customer Lockbox request, the audit record is logged in the audit log contains the following information.

Expand table

Audit record property	Description
Date	The date and time when the Customer Lockbox request was approved or denied.
IP address	The IP address of the machine the approver used to approve or deny a request.
User	The service account BOXServiceAccount@[customerforest].prod.outlook.com.
Activity	Set-AccessToCustomerDataRequest; this is the auditing activity that is logged when you approve or deny a Customer Lockbox request.
Item	The Guid of the Customer Lockbox request

The following screenshot shows an example of an audit record that corresponds to an approved Customer Lockbox request. If a Customer Lockbox request was denied, then the value of `ApprovalDecision` parameter would be `Deny`.



Audit record for an action performed by a Microsoft engineer

The actions performed by a Microsoft engineer after a Customer Lockbox request is approved (and that may result in accessing customer content) are logged in the audit

log. These records contain the following information.

Expand table

Audit record property	Description
Date	Date time when the action was performed. The time that this action was performed will be within 4 hours of when the Customer Lockbox request was approved.
IP address	The IP Address of the machine Microsoft engineer used.
User	Microsoft Operator; this value indicates the record is related to a Customer Lockbox request.
Activity	Name of the activity performed by the Microsoft engineer.
Item	<empty>

Frequently asked questions

Which Microsoft 365 services does Customer Lockbox apply to?

Customer Lockbox is currently supported in Exchange Online, SharePoint Online, OneDrive for Business, and Teams.

Is Customer Lockbox available to all customers?

Customer Lockbox is included with the Microsoft 365 or Office 365 E5 subscriptions and can be added to other plans with an Information Protection and Compliance or an Advanced Compliance add-on subscription. See Plans and pricing for more information.

What is customer content?

Customer content is the data created by users of Microsoft 365 services and applications. Examples of customer content include:

- Email body or email attachments
- SharePoint site contents

- Information in the body of a SharePoint file
- Skype for Business presentation file body
- Instant messages (IM) or voice conversations
- Text entered in Teams chats and Teams channels, for example, 1:1 chats, group chats, shared channels, private channels, and meeting chat
- Other data pasted into Teams chat threads, such as code snippets, images, audio and video messages, and links
- App and bot data in Teams chats and Teams channels
- Teams activity feed
- Teams meeting recordings and transcripts
- Voicemail
- Files posted to Teams chats and Teams channels
- Microsoft Copilot for Microsoft 365 interactions
- Customer-generated blob or structured storage data (for example, SQL Containers)
- Customer-owned security information (for example, certificates, encryption keys, and passwords)
- Inferences, and all subsequent inferences, if customer content remains

For more information about customer content in Office 365, see the Office 365 Trust Center.

Who is notified when there is a request to access my content?

Global administrators and anyone assigned the Customer Lockbox access approver admin role are notified. These are also the same users who can approve for Customer Lockbox requests.

Who can approve or reject these requests in my organization?

Global administrators and anyone assigned the Customer Lockbox access approver admin role can approve Customer Lockbox requests. Customers control these role assignments in their organizations.

How do I opt in to Customer Lockbox?

A global administrator can enable and configure Customer Lockbox in the Microsoft 365 admin center.

If I approve a Customer Lockbox request, what can the engineer do and how will I know what the Microsoft engineer did?

After you approve a Customer Lockbox request, the Microsoft engineer granted these necessary privileges to access customer content by using pre-approved cmdlets. Actions taken by Microsoft engineers in response to Customer Lockbox requests are logged and accessible in the audit log in the Security & Compliance Center.

How do I know that Microsoft follows the approval process?

You can cross-reference the email approval notifications sent to admins and approvers in your organization with the Customer Lockbox request history in the Microsoft 365 admin center.

Customer Lockbox is included in the latest SOC 1 SSAE 16 audit report. For more details, you can find the latest reports in the Microsoft Service Trust Portal.

Can Microsoft modify the list of approvers for my tenant? If not, how is it prevented?

Only a global administrator in your organization can specify who can approve Customer Lockbox requests. That means only the members of the Global administrator group in Microsoft Entra ID can specify who can approve request. Membership of the Global administrator group in Microsoft Entra ID is managed only by your organization.

What if I need more information about a content access request to approve it?

Each Customer Lockbox request contains a Microsoft 365 service request number. You can contact Microsoft Support and reference this service number to get more information about the request.

When a Customer Lockbox request is approved, how long are the permissions valid?

Currently, the maximum period for the access permissions granted to the Microsoft engineer is 4 hours. The Microsoft engineer can also request a shorter period.

How can I get a history of all Customer Lockbox requests?

All Customer Lockbox requests are viewed in the Microsoft 365 admin center.

How do I correlate the content access requests with the related audit logs?

The Compliance Center Activity Feed contains log activities of Customer Lockbox. Customers can cross-reference the Customer Lockbox log activities from the activity feed against the email request they receive.

What happens when a customer doesn't respond to a Customer Lockbox request?

Customer Lockbox requests have a default duration of 12 hours. If you don't respond to a request within 12 hour, the request expires.

What does Microsoft do when a customer rejects a Customer Lockbox request?

If a customer rejects a Customer Lockbox request, no access to customer content occurs. If a user in your organization continues to experience a service issue requiring Microsoft to access customer content to resolve the issue, then the service issue might persist and Microsoft will inform the user about this.

How do I set up alerts whenever a request has been approved?

There is no built-in option to alert administrators. However, administrators can set up alerts using Microsoft Defender for Cloud Apps.

Does Customer Lockbox protect against data requests from law enforcement agencies or other third parties?

No. Microsoft takes third-party requests for customer data seriously. As a cloud service provider, Microsoft always advocates for the privacy of customer data. In the event we get a subpoena, Microsoft always attempts to redirect the third party to the customer to obtain the information. (Read Brad Smith's blog: Protecting customer data from government snooping). We periodically publish detailed information about the law enforcement requests that Microsoft receives.

See the Microsoft Trust Center regarding third-party data requests and the "Disclosure of Customer Data" section in the Online Services Terms for more information.

How does Microsoft ensure that a member of its staff doesn't have standing access to customer content in Office 365 applications?

Microsoft implements extensive preventive measures through access control systems, and detective measures to identify and address attempts to circumvent these access control systems. Microsoft 365 operates with the principles of least privilege and just-in-time access. Therefore, no Microsoft personnel have permission to access customer content on an ongoing basis. If permission is granted, it is for a limited duration.

Microsoft 365 uses an access control system called *Lockbox* to process requests for permissions that grant the ability to perform operational and administrative functions within the service. An operator must request access to customer content using Lockbox, which then requires a second person to take action on the request (for example, approve it) before access is granted. That second person can't be the requestor and must be designated to approve access to customer content. Only if the request is approved does the operator acquire temporary access to customer content. After the elevation period expires, Lockbox revokes access.

Refer to the Online Services Terms for more details about Microsoft general security practices.

Under what circumstances do Microsoft engineers need access to my content?

The most common scenario where Microsoft engineers need access customer content is when the customer makes a support request that requires access for troubleshooting. A foundational principle of Microsoft 365 is that the service operates without Microsoft access to customer content. Nearly all service operations performed by Microsoft are fully automated and human involvement is highly controlled and abstracted away from customer content. The goal for Microsoft 365 is access to customer content to support the service isn't needed until the customer approves a specific request for Microsoft access.

I already thought my data was secure with the Microsoft cloud, so why do I need Customer Lockbox?

Customer Lockbox provides an extra layer of control by offering customers the ability to give explicit access authorization for service operations. By demonstrating that procedures are in place for explicit data access authorization, Customer Lockbox also helps customers meet certain compliance obligations such as HIPAA and FEDRAMP.

Feedback

Was this page helpful? Yes No

Data exfiltration protection access controls

Article • 03/02/2023

Data exfiltration is an attack whereby an internal or external actor completes an unauthorized data transfer of sensitive corporate resources. The exfiltration of sensitive corporate resources is often accomplished due to a lack of appropriate authentication and authorization controls. Microsoft seeks to guard against malicious access and the exfiltration of data to locations outside of their intended organizational scope by employing a suite of mitigating controls to address multiple risk scenarios. By doing so, Microsoft supports the defense-in-depth posture of its customers and reduces the threat of data exfiltration.

This article provides an in-depth look at the holistic approach Microsoft takes for data exfiltration protection on behalf of its customers. It shows how Microsoft's suite of mitigating controls detect and prevent malicious behavior, manage access between tenants, and protect against replay attacks.

Microsoft Defender for Cloud Apps

Actions that would compromise the security of customer data must be detected and prevented. For example, employees may be using an unapproved cloud application for storing sensitive corporate data or downloading a vast number of sensitive files for exfiltration. These actions can be prevented by Microsoft Defender for Cloud Apps.

Microsoft Defender for Cloud Apps is a Cloud Access Security Broker (CASB). CASBs control access in real time between enterprise users and the cloud resources they use, wherever users are located and regardless of the device they are using. Microsoft Defender for Cloud Apps prevents malicious activity across all Microsoft and third-party cloud services by discovering and providing visibility into shadow IT and app use, monitoring user activities for anomalous behavior, controlling access to resources, providing the ability to classify and prevent sensitive information leaks, and assessing the compliance of cloud services. Microsoft Defender for Cloud Apps supports various deployment modes including log collection, API connectors, and reverse proxy. This CASB offering allows organizations to sanction and block the use of specified cloud applications and to block malicious file downloads using Conditional Access App Control.

Microsoft Defender for Cloud Apps Conditional Access App Control (CAAC) uses reverse proxy architecture to provide the tools needed for real-time visibility and control over

access to and activities performed within a cloud environment. CAAC can be used to:

- Avoid data leaks by blocking downloads
- Set rules that force data stored in the cloud to be protected with encryption.
- Highlight unprotected endpoints so that organizations can monitor what's being done on unmanaged devices
- Control access from non-corporate networks or risky IP addresses.
- Reevaluate conditional access policies when a sensitive action occurs in the session. Allowing, for example, the download of a highly confidential file to require multi-factor authentication.

Continuous Access Evaluation

Risky user behavior that precedes data exfiltration can also be detected and prevented, by Continuous Access Evaluation (CAE). CAE works to enforce access policies in near-real time based on signals from users, sessions, and devices, and can revoke sessions when risk is present. For example:

- When a user has been terminated and their access to corporate resources must be immediately revoked.
- When a user authenticates within the corporate boundary, only to walk across the street to a coffee shop, connecting to the riskier network while still being authenticated into cloud resources containing sensitive data.
- When a user account is deleted or disabled.
- When the password for a user is changed or reset.
- When the administrator explicitly revokes all refresh tokens for the user.
- When high user risk is detected by Microsoft Entra ID Protection (device non-compliance, impossible travel, etc.)

CAE can also be used to prevent token export, a common data exfiltration risk. Tokens e.g., access tokens, are provided to an authorized user when a session is initiated. Using developer tools, these tokens can be exported outside of the internal network to external users who then supply the token to a Microsoft 365 service, circumventing Microsoft Entra ID, and gaining access to that resource and the data within. CAE prevents token export by restricting access to corporate resources to trusted IP addresses and always-on VPN. In a token export scenario, CAE would detect the access

request coming from an unauthorized location and revoke access to the resource, preventing data exfiltration.

Mitigating the risk of unauthorized tenant access

Microsoft understands that a key part of business involves collaboration and communication with other enterprises. Collaboration and guest access carry their own inherent risks. Employees can be invited to use their home identities to become guests in another tenant or invite guests from another tenant into their own. Employees might also attempt to use the enterprise accounts of another tenant to access said tenant from within their network, creating an exfiltration channel. Malicious or negligent users may also use their personal Microsoft resources for storage of sensitive corporate data (e.g., Outlook, personal OneDrive). This is why Microsoft employs Tenant Restrictions v2 (TRv2). TRv2 is a cloud-policy based authorization control plane that allows control over employee access to external tenants. With TRv2, users on organizationally managed devices or devices inside the corporate network can be blocked from accessing unsanctioned foreign tenants. Cross-Tenant access policies (XTAP) further enable control over how you collaborate within other organization's tenants (outbound) and how other organizations collaborate within your own organization (inbound). Using XTAP, users can become guests in only explicitly approved tenants.

TRv2 can also be used to prevent token import. Token import takes place when a user logs into an unapproved external tenant from outside the corporate network, then their token is sent to an associate inside the corporate network. The user inside the corporate network then attempts to use the access token to log into an unapproved tenant for the purpose of data exfiltration. TRv2 prevents token import by blocking the internal user from accessing non-allowlisted foreign tenants. In this scenario, the user would not be able to authenticate to the external tenant with the token because of TRv2.

Network security overview

Article • 04/04/2023

How do Microsoft online services secure the network boundary?

Microsoft online services employ multiple strategies for securing its network boundary, including automated detection and prevention of network-based attacks, specialized firewall devices, and Exchange Online Protection (EOP) for anti-spam and anti-malware protection. In addition, Microsoft online services separate their production environments into logically isolated network segments, with only necessary communication permitted between segments. Network traffic is secured using additional network firewalls at boundary points to help detect, prevent, and mitigate network attacks.

How do Microsoft online services defend against DDoS attacks?

Microsoft's large internet presence insulates it from the negative effects of many distributed denial-of-service (DDoS) attacks. Distributed instances of each Microsoft online service and multiple routes to each service limit the impact of DDoS attacks against the system. This redundancy improves Microsoft online services' ability to absorb DDoS attacks and increases the amount of time available to detect and mitigate DDoS attacks before they impact service availability.

In addition to Microsoft's redundant system architecture, Microsoft uses sophisticated detection and mitigation tools to respond to DDoS attacks. Special-purpose firewalls monitor and drop unwanted traffic before it crosses the boundary into the network, reducing stress on systems located inside the network boundary. To further protect our cloud services, Microsoft utilizes a DDoS defense system deployed as part of Microsoft Azure. The Azure DDoS defense system is designed to withstand attacks from the outside and from other Azure tenants.

How does Microsoft protect users against spam and malware being uploaded or sent through online services?

Microsoft online services build antimalware protection into services that might be vectors for malicious code, such as Exchange Online and SharePoint Online. Exchange Online Protection (EOP) scans all emails and email attachments for malware as they enter and exit the system, preventing infected messages and attachments from being delivered. Advanced spam filtering is automatically applied to inbound and outbound messages to help prevent customer organizations from receiving and sending spam. This layer of protection guards against attacks that take advantage of unsolicited or unauthorized email, such as phishing attacks. SharePoint Online uses the same virus detection engine to selectively scan uploaded files for malware. If a file is marked as infected, users are prevented from downloading or syncing the file to protect client endpoints. Similarly, Azure compares hashes related to files uploaded to Azure Storage to those hashes of known malware. When matches are found, an alert is raised in Microsoft Defender for Cloud where a decision is made about the legitimacy of the alert and how it should be addressed.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to network security.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
SOC 1 SOC 2 SOC 3	VM-1: Security event logging VM-3: Intrusion detection and monitoring VM-4: Malicious events investigation VM-6: Vulnerability scanning VM-7: Network device configuration VM-8: Penetration testing VM-9: Network device security event logging VM-13: Network device vulnerability mitigation	November 17, 2023

Microsoft 365

Expand table

External audits	Section	Latest report date
FedRAMP (Office 365)	SC-5: Denial of service protection SC-7: Boundary protection SI-2: Flaw remediation SI-3: Malicious code protection SI-8: Spam protection	July 31, 2023
SOC 1	CA-27: Vulnerability Scanning	January 3, 2023
SOC 2	CA-27: Vulnerability Scanning CA-45: Anti-malware	January 3, 2023

Feedback

Was this page helpful? Yes No

Microsoft denial-of-service defense strategy

Article • 03/02/2023

Denial-of-service defense strategy

Microsoft's strategy to defend against network-based distributed denial-of-service (DDoS) attacks is unique due to a large global footprint, allowing Microsoft to utilize strategies and techniques that are unavailable to most other organizations. Additionally, Microsoft contributes to and draws from collective knowledge aggregated by an extensive threat intelligence network, which includes Microsoft partners and the broader internet security community. This intelligence, along with information gathered from online services and Microsoft's global customer base, continuously improves Microsoft's DDoS defense system that protects all of Microsoft online services' assets.

The cornerstone of Microsoft's DDoS strategy is global presence. Microsoft engages with Internet providers, peering providers (public and private), and private corporations all over the world. This engagement gives Microsoft a significant Internet presence and enables Microsoft to absorb attacks across a large surface area

As Microsoft's edge capacity has grown over time, the significance of attacks against individual edges has substantially diminished. Because of this decrease, Microsoft has separated the detection and mitigation components of its DDoS prevention system. Microsoft deploys multi-tiered detection systems at regional datacenters to detect attacks closer to their saturation points while maintaining global mitigation at the edge nodes. This strategy ensures that Microsoft services can handle multiple simultaneous attacks.

One of the most effective and low-cost defenses employed by Microsoft against DDoS attacks is reducing service attack surfaces. Unwanted traffic is dropped at the network edge instead of analyzing, processing, and scrubbing the data inline.

At the interface with the public network, Microsoft uses special-purpose security devices for firewall, network address translation, and IP filtering functions. Microsoft also uses global equal-cost multi-path (ECMP) routing. Global ECMP routing is a network framework to ensure that there are multiple global paths to reach a service. With multiple paths to each service, DDoS attacks are limited to the region from which the attack originates. Other regions should be unaffected by the attack, as end users would use other paths to reach the service in those regions. Microsoft has also developed

internal DDoS correlation and detection systems that use flow data, performance metrics, and other information to rapidly detect DDoS attacks.

To further protect cloud services, Microsoft uses Azure DDoS Protection, a DDoS defense system built into Microsoft Azure's continuous monitoring and penetration-testing processes. Azure DDoS Protection is designed not only to withstand external attacks, but also attacks from other Azure tenants. Azure uses standard detection and mitigation techniques such as SYN cookies, rate limiting, and connection limits to protect against DDoS attacks. To support automated protections, a cross-workload DDoS incident response team identifies the roles and responsibilities across teams, the criteria for escalations, and the protocols for incident handling across affected teams.

Most DDoS attacks launched against targets are at the Network (L3) and Transport (L4) layers of the Open Systems Interconnection (OSI) model. Attacks directed at the L3 and L4 layers are designed to flood a network interface or service with attack traffic to overwhelm resources and deny the ability to respond to legitimate traffic. To guard against L3 and L4 attacks, Microsoft's DDoS solutions use traffic sampling data from datacenter routers to safeguard the infrastructure and customer targets. Traffic sampling data is analyzed by a network monitoring service to detect attacks. When an attack is detected, automated defense mechanisms kick in to mitigate the attack and ensure that attack traffic directed at one customer does not result in collateral damage or diminished network quality of service for other customers.

Microsoft also takes an offensive approach to DDoS defense. Botnets are a common source of command and control for conducting DDoS attacks to amplify attacks and maintain anonymity. The Microsoft Digital Crimes Unit (DCU) focuses on identifying, investigating, and disrupting malware distribution and communications infrastructure to reduce the scale and impact of botnets.

Application-level defenses

Microsoft's cloud services are intentionally built to support high loads, which help to protect against application-level DDoS attacks. Microsoft's scaled-out architecture distributes services across multiple global datacenters with regional isolation and workload-specific throttling features for relevant workloads.

Each customer's country or region, which the customer's administrator identifies during the initial configuration of the services, determines the primary storage location for that customer's data. Customer data is replicated between redundant datacenters according to a primary/backup strategy. A primary datacenter hosts the application software along with all the primary customer data running on the software. A backup datacenter provides automatic failover. If the primary datacenter ceases to function for any reason,

requests are redirected to the copy of the software and customer data in the backup datacenter. At any given time, customer data may be processed in either the primary or the backup datacenter. Distributing data across multiple datacenters reduces the affected surface area in case one datacenter is attacked. Furthermore, the services in the affected datacenter can be quickly redirected to the secondary datacenter to maintain availability during an attack and redirected back to the primary datacenter once an attack has been mitigated.

As another mitigation against DDoS attacks, individual workloads include built-in features that manage resource utilization. For example, the throttling mechanisms in Exchange Online and SharePoint Online are part of a multi-layered approach to defending against DDoS attacks.

Azure SQL Database has an extra layer of security in the form of a gateway service called DoSGuard that tracks failed login attempts based on IP address. If the threshold for failed login attempts from the same IP is reached, DoSGuard blocks the address for a pre-determined amount of time.

Resources

- Azure DDoS Protection Standard overview
- Azure DDoS Protection Standard fundamental best practices
- Components of a DDoS response strategy

The Azure production network

Article • 04/02/2023

The users of the Azure production network include both external customers who access their own Azure applications and internal Azure support personnel who manage the production network. This article discusses the security access methods and protection mechanisms for establishing connections to the Azure production network.

Internet routing and fault tolerance

A globally redundant internal and external Azure Domain Name Service (DNS) infrastructure, combined with multiple primary and secondary DNS server clusters, provides fault tolerance. At the same time, additional Azure network security controls, such as NetScaler, are used to prevent distributed denial of service (DDoS) attacks and protect the integrity of Azure DNS services.

The Azure DNS servers are located at multiple datacenter facilities. The Azure DNS implementation incorporates a hierarchy of secondary and primary DNS servers to publicly resolve Azure customer domain names. The domain names usually resolve to a CloudApp.net address, which wraps the virtual IP (VIP) address for the customer's service. Unique to Azure, the VIP that corresponds to internal dedicated IP (DIP) address of the tenant translation is done by the Microsoft load balancers responsible for that VIP.

Azure is hosted in geographically distributed Azure datacenters within the US, and it's built on state-of-the-art routing platforms that implement robust, scalable architectural standards. Among the notable features are:

- Multiprotocol Label Switching (MPLS)-based traffic engineering, which provides efficient link utilization and graceful degradation of service if there is an outage.
- Networks are implemented with "need plus one" (N+1) redundancy architectures or better.
- Externally, datacenters are served by dedicated, high-bandwidth network circuits that redundantly connect properties with over 1,200 internet service providers globally at multiple peering points. This connection provides in excess of 2,000 gigabytes per second (GBps) of edge capacity.

Because Microsoft owns its own network circuits between datacenters, these attributes help the Azure offering achieve 99.9+ percent network availability without the need for traditional third-party internet service providers.

Connection to production network and associated firewalls

The Azure network internet traffic flow policy directs traffic to the Azure production network that's located in the nearest regional datacenter within the US. Because the Azure production datacenters maintain consistent network architecture and hardware, the traffic flow description that follows applies consistently to all datacenters.

After internet traffic for Azure is routed to the nearest datacenter, a connection is established to the access routers. These access routers serve to isolate traffic between Azure nodes and customer-instantiated VMs. Network infrastructure devices at the access and edge locations are the boundary points where ingress and egress filters are applied. These routers are configured through a tiered access-control list (ACL) to filter unwanted network traffic and apply traffic rate limits, if necessary. Traffic that is allowed by ACL is routed to the load balancers. Distribution routers are designed to allow only Microsoft-approved IP addresses, provide anti-spoofing, and establish TCP connections that use ACLs.

External load-balancing devices are located behind the access routers to perform network address translation (NAT) from internet-routable IPs to Azure internal IPs. The devices also route packets to valid production internal IPs and ports, and they act as a protection mechanism to limit exposing the internal production network address space.

By default, Microsoft enforces Hypertext Transfer Protocol Secure (HTTPS) for all traffic that's transmitted to customers' web browsers, including sign-in and all traffic thereafter. The use of TLS v1.2 enables a secure tunnel for traffic to flow through. ACLs on access and core routers ensure that the source of the traffic is consistent with what is expected.

An important distinction in this architecture, when it's compared to traditional security architecture, is that there are no dedicated hardware firewalls, specialized intrusion detection or prevention devices, or other security appliances that are normally expected before connections are made to the Azure production environment. Customers usually expect these hardware firewall devices in the Azure network; however, none are employed within Azure. Almost exclusively, those security features are built into the software that runs the Azure environment to provide robust, multi-layered security mechanisms, including firewall capabilities. Additionally, the scope of the boundary and associated sprawl of critical security devices is easier to manage and inventory, as shown in the preceding illustration, because it is managed by the software that's running Azure.

Core security and firewall features

Azure implements robust software security and firewall features at various levels to enforce security features that are usually expected in a traditional environment to protect the core Security Authorization boundary.

Azure security features

Azure implements host-based software firewalls inside the production network. Several core security and firewall features reside within the core Azure environment. These security features reflect a defense-in-depth strategy within the Azure environment. Customer data in Azure is protected by the following firewalls:

Hypervisor firewall (packet filter): This firewall is implemented in the hypervisor and configured by the fabric controller (FC) agent. This firewall protects the tenant that runs inside the VM from unauthorized access. By default, when a VM is created, all traffic is blocked and then the FC agent adds rules and exceptions in the filter to allow authorized traffic.

Two categories of rules are programmed here:

- **Machine config or infrastructure rules**: By default, all communication is blocked. Exceptions exist that allow a VM to send and receive Dynamic Host Configuration Protocol (DHCP) communications and DNS information, and send traffic to the "public" internet outbound to other VMs within the FC cluster and OS Activation server. Because the VMs' allowed list of outgoing destinations does not include Azure router subnets and other Microsoft properties, the rules act as a layer of defense for them.
- **Role configuration file rules**: Defines the inbound ACLs based on the tenants' service model. For example, if a tenant has a web front end on port 80 on a certain VM, port 80 is opened to all IP addresses. If the VM has a worker role running, the worker role is opened only to the VM within the same tenant.

Native host firewall: Azure Service Fabric and Azure Storage run on a native OS, which has no hypervisor and, therefore, Windows Firewall is configured with the preceding two sets of rules.

Host firewall: The host firewall protects the host partition, which runs the hypervisor. The rules are programmed to allow only the FC and jump boxes to talk to the host partition on a specific port. The other exceptions are to allow DHCP response and DNS replies. Azure uses a machine configuration file, which contains a template of firewall rules for the host partition. A host firewall exception also exists that allows VMs to communicate to host components, wire server, and metadata server, through specific protocol/ports.

Guest firewall: The Windows Firewall piece of the guest OS, which is configurable by customers on customer VMs and storage.

Additional security features that are built into the Azure capabilities include:

- Infrastructure components that are assigned IP addresses that are from DIPs. An attacker on the internet cannot address traffic to those addresses because it would not reach Microsoft. Internet gateway routers filter packets that are addressed solely to internal addresses, so they would not enter the production network. The only components that accept traffic that's directed to VIPs are load balancers.

- Firewalls that are implemented on all internal nodes have three primary security architecture considerations for any given scenario:
 - Firewalls are placed behind the load balancer and accept packets from anywhere. These packets are intended to be externally exposed and would correspond to the open ports in a traditional perimeter firewall.
 - Firewalls accept packets only from a limited set of addresses. This consideration is part of the defensive in-depth strategy against DDoS attacks. Such connections are cryptographically authenticated.
 - Firewalls can be accessed only from select internal nodes. They accept packets only from an enumerated list of source IP addresses, all of which are DIPs within the Azure network. For example, an attack on the corporate network could direct requests to these addresses, but the attacks would be blocked unless the source address of the packet was one in the enumerated list within the Azure network.
 - The access router at the perimeter blocks outbound packets that are addressed to an address that's inside the Azure network because of its configured static routes.

Next steps

To learn more about what Microsoft does to secure the Azure infrastructure, see:

- Azure facilities, premises, and physical security
- Azure infrastructure availability
- Azure information system components and boundaries
- Azure network architecture
- Azure SQL Database security features
- Azure production operations and management
- Azure infrastructure monitoring
- Azure infrastructure integrity
- Azure customer data protection

Azure network security overview

Article • 04/02/2023

Network security could be defined as the process of protecting resources from unauthorized access or attack by applying controls to network traffic. The goal is to ensure that only legitimate traffic is allowed. Azure includes a robust networking infrastructure to support your application and service connectivity requirements. Network connectivity is possible between resources located in Azure, between on-premises and Azure hosted resources, and to and from the internet and Azure.

This article covers some of the options that Azure offers in the area of network security. You can learn about:

- Azure networking
- Network access control
- Azure Firewall
- Secure remote access and cross-premises connectivity
- Availability
- Name resolution
- Perimeter network (DMZ) architecture
- Azure DDoS protection
- Azure Front Door
- Traffic manager
- Monitoring and threat detection

> **Note**
>
> For web workloads, we highly recommend utilizing **Azure DDoS protection** and a **web application firewall** to safeguard against emerging DDoS attacks. Another option is to deploy **Azure Front Door** along with a web application firewall. Azure Front Door offers platform-level **protection against network-level DDoS attacks**.

Azure networking

Azure requires virtual machines to be connected to an Azure Virtual Network. A virtual network is a logical construct built on top of the physical Azure network fabric. Each virtual network is isolated from all other virtual networks. This helps ensure that network traffic in your deployments is not accessible to other Azure customers.

Learn more:

- Virtual network overview

Network access control

Network access control is the act of limiting connectivity to and from specific devices or subnets within a virtual network. The goal of network access control is to limit access to your virtual machines and services to approved users and devices. Access controls are based on decisions to allow or deny connections to and from your virtual machine or service.

Azure supports several types of network access control, such as:

- Network layer control
- Route control and forced tunneling
- Virtual network security appliances

Network layer control

Any secure deployment requires some measure of network access control. The goal of network access control is to restrict virtual machine communication to the necessary systems. Other communication attempts are blocked.

> ⓘ **Note**
>
> Storage Firewalls are covered in the **Azure storage security overview** article

Network security rules (NSGs)

If you need basic network level access control (based on IP address and the TCP or UDP protocols), you can use Network Security Groups (NSGs). An NSG is a basic, stateful, packet filtering firewall, and it enables you to control access based on a 5-tuple. NSGs include functionality to simplify management and reduce the chances of configuration mistakes:

- **Augmented security rules** simplify NSG rule definition and allow you to create complex rules rather than having to create multiple simple rules to achieve the same result.
- **Service tags** are Microsoft created labels that represent a group of IP addresses. They update dynamically to include IP ranges that meet the conditions that define inclusion in the label. For example, if you want to create a rule that applies to all Azure storage on the east region you can use Storage.EastUS

- **Application security groups** allow you to deploy resources to application groups and control the access to those resources by creating rules that use those application groups. For example, if you have webservers deployed to the 'Webservers' application group you can create a rule that applies a NSG allowing 443 traffic from the Internet to all systems in the 'Webservers' application group.

NSGs do not provide application layer inspection or authenticated access controls.

Learn more:

- Network Security Groups

Defender for Cloud just in time VM access

Microsoft Defender for Cloud can manage the NSGs on VMs and lock access to the VM until a user with the appropriate Azure role-based access control Azure RBAC permissions requests access. When the user is successfully authorized Defender for Cloud makes modifications to the NSGs to allow access to selected ports for the time specified. When the time expires the NSGs are restored to their previous secured state.

Learn more:

- Microsoft Defender for Cloud Just in Time Access

Service endpoints

Service endpoints are another way to apply control over your traffic. You can limit communication with supported services to just your VNets over a direct connection. Traffic from your VNet to the specified Azure service remains on the Microsoft Azure backbone network.

Learn more:

- Service endpoints

Route control and forced tunneling

The ability to control routing behavior on your virtual networks is critical. If routing is configured incorrectly, applications and services hosted on your virtual machine might connect to unauthorized devices, including systems owned and operated by potential attackers.

Azure networking supports the ability to customize the routing behavior for network traffic on your virtual networks. This enables you to alter the default routing table entries in your virtual network. Control of routing behavior helps you make sure that all traffic from a certain device or group of devices enters or leaves your virtual network through a specific location.

For example, you might have a virtual network security appliance on your virtual network. You want to make sure that all traffic to and from your virtual network goes through that virtual security appliance. You can do this by configuring User Defined Routes (UDRs) in Azure.

Forced tunneling is a mechanism you can use to ensure that your services are not allowed to initiate a connection to devices on the internet. Note that this is different from accepting incoming connections and then responding to them. Front-end web servers need to respond to requests from internet hosts, and so internet-sourced traffic is allowed inbound to these web servers and the web servers are allowed to respond.

What you don't want to allow is a front-end web server to initiate an outbound request. Such requests might represent a security risk because these connections can be used to download malware. Even if you do want these front-end servers to initiate outbound requests to the internet, you might want to force them to go through your on-premises web proxies. This enables you to take advantage of URL filtering and logging.

Instead, you would want to use forced tunneling to prevent this. When you enable forced tunneling, all connections to the internet are forced through your on-premises gateway. You can configure forced tunneling by taking advantage of UDRs.

Learn more:

- What are User Defined Routes and IP Forwarding

Virtual network security appliances

While NSGs, UDRs, and forced tunneling provide you a level of security at the network and transport layers of the OSI model, you might also want to enable security at levels higher than the network.

For example, your security requirements might include:

- Authentication and authorization before allowing access to your application
- Intrusion detection and intrusion response
- Application layer inspection for high-level protocols
- URL filtering

- Network level antivirus and Antimalware
- Anti-bot protection
- Application access control
- Additional DDoS protection (above the DDoS protection provided by the Azure fabric itself)

You can access these enhanced network security features by using an Azure partner solution. You can find the most current Azure partner network security solutions by visiting the Azure Marketplace, and searching for "security" and "network security."

Azure Firewall

Azure Firewall is a cloud-native and intelligent network firewall security service that provides threat protection for your cloud workloads running in Azure. It's a fully stateful firewall as a service with built-in high availability and unrestricted cloud scalability. It provides both east-west and north-south traffic inspection.

Azure Firewall is offered in three SKUs: Standard, Premium, and Basic. Azure Firewall Standard provides L3-L7 filtering and threat intelligence feeds directly from Microsoft Cyber Security. Azure Firewall Premium provides advanced capabilities include signature-based IDPS to allow rapid detection of attacks by looking for specific patterns. Azure Firewall Basic is a simplified SKU that provides the same level of security as the Standard SKU but without the advanced capabilities.

Learn more:

- What is Azure Firewall

Secure remote access and cross-premises connectivity

Setup, configuration, and management of your Azure resources needs to be done remotely. In addition, you might want to deploy hybrid IT solutions that have components on-premises and in the Azure public cloud. These scenarios require secure remote access.

Azure networking supports the following secure remote access scenarios:

- Connect individual workstations to a virtual network
- Connect your on-premises network to a virtual network with a VPN
- Connect your on-premises network to a virtual network with a dedicated WAN link
- Connect virtual networks to each other

Connect individual workstations to a virtual network

You might want to enable individual developers or operations personnel to manage virtual machines and services in Azure. For example, let's say you need access to a virtual machine on a virtual network. But your security policy does not allow RDP or SSH remote access to individual virtual machines. In this case, you can use a point-to-site VPN connection.

The point-to-site VPN connection enables you to set up a private and secure connection between the user and the virtual network. When the VPN connection is established, the user can RDP or SSH over the VPN link into any virtual machine on the virtual network. (This assumes that the user can authenticate and is authorized.) Point-to-site VPN supports:

- Secure Socket Tunneling Protocol (SSTP), a proprietary SSL-based VPN protocol. An SSL VPN solution can penetrate firewalls, since most firewalls open TCP port 443, which TLS/SSL uses. SSTP is only supported on Windows devices. Azure supports all versions of Windows that have SSTP (Windows 7 and later).
- IKEv2 VPN, a standards-based IPsec VPN solution. IKEv2 VPN can be used to connect from Mac devices (OSX versions 10.11 and above).
- OpenVPN

Learn more:

- Configure a point-to-site connection to a virtual network using PowerShell

Connect your on-premises network to a virtual network with a VPN

You might want to connect your entire corporate network, or portions of it, to a virtual network. This is common in hybrid IT scenarios, where organizations extend their on-premises datacenter into Azure. In many cases, organizations host parts of a service in Azure, and parts on-premises. For example,they might do so when a solution includes front-end web servers in Azure and back-end databases on-premises. These types of "cross-premises" connections also make management of Azure located resources more secure, and enable scenarios such as extending Active Directory domain controllers into Azure.

One way to accomplish this is to use a site-to-site VPN. The difference between a site-to-site VPN and a point-to-site VPN is that the latter connects a single device to a virtual network. A site-to-site VPN connects an entire network (such as your on-premises

network) to a virtual network. Site-to-site VPNs to a virtual network use the highly secure IPsec tunnel mode VPN protocol.

Learn more:

- Create a Resource Manager VNet with a site-to-site VPN connection using the Azure portal
- About VPN Gateway

Connect your on-premises network to a virtual network with a dedicated WAN link

Point-to-site and site-to-site VPN connections are effective for enabling cross-premises connectivity. However, some organizations consider them to have the following drawbacks:

- VPN connections move data over the internet. This exposes these connections to potential security issues involved with moving data over a public network. In addition, reliability and availability for internet connections cannot be guaranteed.
- VPN connections to virtual networks might not have the bandwidth for some applications and purposes, as they max out at around 200 Mbps.

Organizations that need the highest level of security and availability for their cross-premises connections typically use dedicated WAN links to connect to remote sites. Azure provides you the ability to use a dedicated WAN link that you can use to connect your on-premises network to a virtual network. Azure ExpressRoute, Express route direct, and Express route global reach enable this.

Learn more:

- ExpressRoute technical overview
- ExpressRoute direct
- Express route global reach

Connect virtual networks to each other

It is possible to use many virtual networks for your deployments. There are various reasons why you might do this. You might want to simplify management, or you might want increased security. Regardless of the motivation for putting resources on different virtual networks, there might be times when you want resources on each of the networks to connect with one another.

One option is for services on one virtual network to connect to services on another virtual network, by "looping back" through the internet. The connection starts on one virtual network, goes through the internet, and then comes back to the destination virtual network. This option exposes the connection to the security issues inherent in any internet-based communication.

A better option might be to create a site-to-site VPN that connects between two virtual networks. This method uses the same IPSec tunnel mode protocol as the cross-premises site-to-site VPN connection mentioned above.

The advantage of this approach is that the VPN connection is established over the Azure network fabric, instead of connecting over the internet. This provides you an extra layer of security, compared to site-to-site VPNs that connect over the internet.

Learn more:

- Configure a VNet-to-VNet Connection by using Azure Resource Manager and PowerShell

Another way to connect your virtual networks is VNET peering. This feature allows you to connect two Azure networks so that communication between them happens over the Microsoft backbone infrastructure without it ever going over the Internet. VNET peering can connect two VNETs within the same region or two VNETs across Azure regions. NSGs can be used to limit connectivity between different subnets or systems.

Availability

Availability is a key component of any security program. If your users and systems can't access what they need to access over the network, the service can be considered compromised. Azure has networking technologies that support the following high-availability mechanisms:

- HTTP-based load balancing
- Network level load balancing
- Global load balancing

Load balancing is a mechanism designed to equally distribute connections among multiple devices. The goals of load balancing are:

- To increase availability. When you load balance connections across multiple devices, one or more of the devices can become unavailable without compromising the service. The services running on the remaining online devices can continue to serve the content from the service.

- To increase performance. When you load balance connections across multiple devices, a single device doesn't have to handle all processing. Instead, the processing and memory demands for serving the content is spread across multiple devices.

HTTP-based load balancing

Organizations that run web-based services often desire to have an HTTP-based load balancer in front of those web services. This helps ensure adequate levels of performance and high availability. Traditional, network-based load balancers rely on network and transport layer protocols. HTTP-based load balancers, on the other hand, make decisions based on characteristics of the HTTP protocol.

Azure Application Gateway provides HTTP-based load balancing for your web-based services. Application Gateway supports:

- Cookie-based session affinity. This capability makes sure that connections established to one of the servers behind that load balancer stays intact between the client and server. This ensures stability of transactions.
- TLS offload. When a client connects with the load balancer, that session is encrypted by using the HTTPS (TLS) protocol. However, in order to increase performance, you can use the HTTP (unencrypted) protocol to connect between the load balancer and the web server behind the load balancer. This is referred to as "TLS offload," because the web servers behind the load balancer don't experience the processor overhead involved with encryption. The web servers can therefore service requests more quickly.
- URL-based content routing. This feature makes it possible for the load balancer to make decisions about where to forward connections based on the target URL. This provides a lot more flexibility than solutions that make load balancing decisions based on IP addresses.

Learn more:

- Application Gateway overview

Network level load balancing

In contrast to HTTP-based load balancing, network level load balancing makes decisions based on IP address and port (TCP or UDP) numbers. You can gain the benefits of network level load balancing in Azure by using Azure Load Balancer. Some key characteristics of Load Balancer include:

- Network level load balancing based on IP address and port numbers.
- Support for any application layer protocol.
- Load balances to Azure virtual machines and cloud services role instances.
- Can be used for both internet-facing (external load balancing) and non-internet facing (internal load balancing) applications and virtual machines.
- Endpoint monitoring, which is used to determine if any of the services behind the load balancer have become unavailable.

Learn more:

- Internal load balancer overview

Global load balancing

Some organizations want the highest level of availability possible. One way to reach this goal is to host applications in globally distributed datacenters. When an application is hosted in datacenters located throughout the world, it's possible for an entire geopolitical region to become unavailable, and still have the application up and running.

This load-balancing strategy can also yield performance benefits. You can direct requests for the service to the datacenter that is nearest to the device that is making the request.

In Azure, you can gain the benefits of global load balancing by using Azure Traffic Manager.

Learn more:

- What is Traffic Manager?

Name resolution

Name resolution is a critical function for all services you host in Azure. From a security perspective, compromise of the name resolution function can lead to an attacker redirecting requests from your sites to an attacker's site. Secure name resolution is a requirement for all your cloud hosted services.

There are two types of name resolution you need to address:

- Internal name resolution. This is used by services on your virtual networks, your on-premises networks, or both. Names used for internal name resolution are not accessible over the internet. For optimal security, it's important that your internal name resolution scheme is not accessible to external users.

- External name resolution. This is used by people and devices outside of your on-premises networks and virtual networks. These are the names that are visible to the internet, and are used to direct connection to your cloud-based services.

For internal name resolution, you have two options:

- A virtual network DNS server. When you create a new virtual network, a DNS server is created for you. This DNS server can resolve the names of the machines located on that virtual network. This DNS server is not configurable, is managed by the Azure fabric manager, and can therefore help you secure your name resolution solution.
- Bring your own DNS server. You have the option of putting a DNS server of your own choosing on your virtual network. This DNS server can be an Active Directory integrated DNS server, or a dedicated DNS server solution provided by an Azure partner, which you can obtain from the Azure Marketplace.

Learn more:

- Virtual network overview
- Manage DNS Servers used by a virtual network

For external name resolution, you have two options:

- Host your own external DNS server on-premises.
- Host your own external DNS server with a service provider.

Many large organizations host their own DNS servers on-premises. They can do this because they have the networking expertise and global presence to do so.

In most cases, it's better to host your DNS name resolution services with a service provider. These service providers have the network expertise and global presence to ensure very high availability for your name resolution services. Availability is essential for DNS services, because if your name resolution services fail, no one will be able to reach your internet facing services.

Azure provides you with a highly available and high-performing external DNS solution in the form of Azure DNS. This external name resolution solution takes advantage of the worldwide Azure DNS infrastructure. It allows you to host your domain in Azure, using the same credentials, APIs, tools, and billing as your other Azure services. As part of Azure, it also inherits the strong security controls built into the platform.

Learn more:

- Azure DNS overview

- Azure DNS private zones allows you to configure private DNS names for Azure resources rather than the automatically assigned names without the need to add a custom DNS solution.

Perimeter network architecture

Many large organizations use perimeter networks to segment their networks, and create a buffer-zone between the internet and their services. The perimeter portion of the network is considered a low-security zone, and no high-value assets are placed in that network segment. You'll typically see network security devices that have a network interface on the perimeter network segment. Another network interface is connected to a network that has virtual machines and services that accept inbound connections from the internet.

You can design perimeter networks in a number of different ways. The decision to deploy a perimeter network, and then what type of perimeter network to use if you decide to use one, depends on your network security requirements.

Learn more:

- Perimeter networks for security zones

Azure DDoS protection

Distributed denial of service (DDoS) attacks are some of the largest availability and security concerns facing customers that are moving their applications to the cloud. A DDoS attack attempts to exhaust an application's resources, making the application unavailable to legitimate users. DDoS attacks can be targeted at any endpoint that is publicly reachable through the internet.

DDoS Protection features include:

- **Native platform integration:** Natively integrated into Azure. Includes configuration through the Azure portal. DDoS Protection understands your resources and resource configuration.
- **Turn-key protection:** Simplified configuration immediately protects all resources on a virtual network as soon as DDoS Protection is enabled. No intervention or user definition is required. DDoS Protection instantly and automatically mitigates the attack, once it is detected.
- **Always-on traffic monitoring:** Your application traffic patterns are monitored 24 hour a day, 7 days a week, looking for indicators of DDoS attacks. Mitigation is performed when protection policies are exceeded.

- **Attack Mitigation Reports** Attack Mitigation Reports use aggregated network flow data to provide detailed information about attacks targeted at your resources.
- **Attack Mitigation Flow Logs** Attack Mitigation Flow Logs allow you to review the dropped traffic, forwarded traffic and other attack data in near real-time during an active DDoS attack.
- **Adaptive tuning:** Intelligent traffic profiling learns your application's traffic over time, and selects and updates the profile that is the most suitable for your service. The profile adjusts as traffic changes over time. Layer 3 to layer 7 protection: Provides full stack DDoS protection, when used with a web application firewall.
- **Extensive mitigation scale:** Over 60 different attack types can be mitigated, with global capacity, to protect against the largest known DDoS attacks.
- **Attack metrics:** Summarized metrics from each attack are accessible through Azure Monitor.
- **Attack alerting:** Alerts can be configured at the start and stop of an attack, and over the attack's duration, using built-in attack metrics. Alerts integrate into your operational software like Microsoft Azure Monitor logs, Splunk, Azure Storage, Email, and the Azure portal.
- **Cost guarantee:** Data-transfer and application scale-out service credits for documented DDoS attacks.
- **DDoS Rapid responsive** DDoS Protection customers now have access to Rapid Response team during an active attack. DRR can help with attack investigation, custom mitigations during an attack and post-attack analysis.

Learn more:

- DDOS protection overview

Azure Front Door

Azure Front Door Service enables you to define, manage, and monitor the global routing of your web traffic. It optimizes your traffic's routing for best performance and high availability. Azure Front Door allows you to author custom web application firewall (WAF) rules for access control to protect your HTTP/HTTPS workload from exploitation based on client IP addresses, country code, and http parameters. Additionally, Front Door also enables you to create rate limiting rules to battle malicious bot traffic, it includes TLS offloading and per-HTTP/HTTPS request, application-layer processing.

Front Door platform itself is protected by an Azure infrastructure-level DDoS protection. For further protection, Azure DDoS Network Protection may be enabled at your VNETs and safeguard resources from network layer (TCP/UDP) attacks via auto tuning and

mitigation. Front Door is a layer 7 reverse proxy, it only allows web traffic to pass through to back end servers and block other types of traffic by default.

> ⓘ **Note**
>
> For web workloads, we highly recommend utilizing **Azure DDoS protection** and a **web application firewall** to safeguard against emerging DDoS attacks. Another option is to deploy **Azure Front Door** along with a web application firewall. Azure Front Door offers platform-level **protection against network-level DDoS attacks**.

Learn more:

- For more information on the whole set of Azure Front door capabilities you can review the Azure Front Door overview

Azure Traffic manager

Azure Traffic Manager is a DNS-based traffic load balancer that enables you to distribute traffic optimally to services across global Azure regions, while providing high availability and responsiveness. Traffic Manager uses DNS to direct client requests to the most appropriate service endpoint based on a traffic-routing method and the health of the endpoints. An endpoint is any Internet-facing service hosted inside or outside of Azure. Traffic manager monitors the end points and does not direct traffic to any endpoints that are unavailable.

Learn more:

- Azure Traffic manager overview

Monitoring and threat detection

Azure provides capabilities to help you in this key area with early detection, monitoring, and collecting and reviewing network traffic.

Azure Network Watcher

Azure Network Watcher can help you troubleshoot, and provides a whole new set of tools to assist with the identification of security issues.

Security Group View helps with auditing and security compliance of Virtual Machines. Use this feature to perform programmatic audits, comparing the baseline policies

defined by your organization to effective rules for each of your VMs. This can help you identify any configuration drift.

Packet capture allows you to capture network traffic to and from the virtual machine. You can collect network statistics and troubleshoot application issues, which can be invaluable in the investigation of network intrusions. You can also use this feature together with Azure Functions to start network captures in response to specific Azure alerts.

For more information on Network Watcher and how to start testing some of the functionality in your labs, see Azure network watcher monitoring overview.

> **Note**
>
> For the most up-to-date notifications on availability and status of this service, check the **Azure updates page**.

Microsoft Defender for Cloud

Microsoft Defender for Cloud helps you prevent, detect, and respond to threats, and provides you increased visibility into, and control over, the security of your Azure resources. It provides integrated security monitoring and policy management across your Azure subscriptions, helps detect threats that might otherwise go unnoticed, and works with a large set of security solutions.

Defender for Cloud helps you optimize and monitor network security by:

- Providing network security recommendations.
- Monitoring the state of your network security configuration.
- Alerting you to network based threats, both at the endpoint and network levels.

Learn more:

- Introduction to Microsoft Defender for Cloud

Virtual Network TAP

Azure virtual network TAP (Terminal Access Point) allows you to continuously stream your virtual machine network traffic to a network packet collector or analytics tool. The collector or analytics tool is provided by a network virtual appliance partner. You can use the same virtual network TAP resource to aggregate traffic from multiple network interfaces in the same or different subscriptions.

Learn more:

- Virtual network TAP

Logging

Logging at a network level is a key function for any network security scenario. In Azure, you can log information obtained for NSGs to get network level logging information. With NSG logging, you get information from:

- Activity logs. Use these logs to view all operations submitted to your Azure subscriptions. These logs are enabled by default, and can be used within the Azure portal. They were previously known as audit or operational logs.
- Event logs. These logs provide information about what NSG rules were applied.
- Counter logs. These logs let you know how many times each NSG rule was applied to deny or allow traffic.

You can also use Microsoft Power BI, a powerful data visualization tool, to view and analyze these logs. Learn more:

- Azure Monitor logs for Network Security Groups (NSGs)

Personnel management overview

Article • 04/04/2023

How does Microsoft screen prospective employees?

Microsoft follows rigorous personnel screening requirements for all candidates, which includes full-time, part-time employees and interns. All candidates are screened prior to beginning employment at Microsoft.

Background checks on employment candidates generally include review of the following components, to the extent permitted by law:

- Identity check
- Education verification
- Employment verification
- Criminal record review
- Sex offender registry review
- Global sanctions list review

What additional checks are performed for employees that manage cloud services?

In addition to pre-employment screening, Microsoft employees who maintain Microsoft online services in the United States must undergo a Microsoft Cloud Background Check as a prerequisite for access to online services systems. The requirements of background check vary to comply with applicable laws and service delivery models. The results from the Microsoft Cloud Background Check are stored in our employee database and must be renewed every two years at a minimum. If the Microsoft Cloud Background Check expires and the employee doesn't renew it, access to online services is revoked and no longer available until the Microsoft Cloud Background Check is completed. Likewise, when the employment relationship with Microsoft ends, all access is immediately revoked.

How does Microsoft ensure employees maintain sufficient skills and knowledge to

perform their responsibilities and follow Microsoft policies?

All Microsoft employees are required to complete basic security awareness training. Initial training occurs when a new employee begins working at Microsoft, and annual refresher training takes place every year thereafter. The training is designed to provide the employee with an understanding of Microsoft's fundamental approach to security. Applicable role-based security training is also required prior to granting any specific access needed for an individual's job responsibilities. Microsoft employees' security training is refreshed on an annual basis, and when system or policy changes warrant new training.

In addition to security awareness training, Microsoft employees must complete Standards of Business Conduct training. This training includes business ethics, employee safety, privacy, anti-harassment, and zero tolerance for non-ethical behavior. At the end of the course, employees must attest that they'll abide by the Microsoft code of business conduct, which is tracked at the organization level. The Standards of Business Conduct training are refreshed on an annual basis.

How does Microsoft revoke access for employees who leave Microsoft?

Microsoft uses clearly defined policies and procedures to promptly revoke physical and logical access to Microsoft systems and resources when an employee leaves Microsoft or is terminated. Microsoft's termination process ensures that former Microsoft employees can't access data or systems after their employment ends.

When a service team user's employment is marked as terminated, this information propagates to the Microsoft account management tool, which automatically removes the terminated employee's domain account. Any access badges or other physical authenticators issued to the terminated employee are collected at the time of the exit interview or termination.

How does Microsoft ensure third-party suppliers meet the same personnel requirements as Microsoft employees?

Microsoft online services require third-party suppliers to have a signed Master Supplier Services Agreement (MSSA). This agreement requires the supplier to comply with

Microsoft policies and procedures, including personnel security policies and procedures. Microsoft monitors compliance with screening requirements for third-party personnel by tracking the outcome of screening directly. Microsoft requires suppliers to conduct background screens for all persons who need access to Microsoft's facilities and/or network. For specific roles, a Supplier may be required to provide attestation as evidence that the person completed the cloud background screen requirements.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to human resources.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
ISO 27001/27002 Statement of Applicability Certificate	A.7: Human resource security	April 24, 2023
ISO 27017 Statement of Applicability Certificate	A.7: Human resource security	April 24, 2023
SOC 1	IS-4: Security training OA-3: Account revocation	November 17, 2023
SOC 2 SOC 3	C5-2: Supplier risk assessment ELC-6: Supplier code of conduct IS-4: Security training OA-3: Account revocation SOC2-1: Disciplinary actions SOC2-12: Background checks SOC2-13: Employment agreements SOC2-14: Confidentiality and non-disclosure agreements	November 17, 2023

Microsoft 365

Expand table

External audits	Section	Latest report date
FedRAMP (Office 365)	AT-2: Security awareness AT-3: Role-based security training AT-4: Security training records PS-3: Personnel screening PS-4: Personnel termination PS-5: Personnel transfer PS-7: Third-party personnel security	July 31, 2023
ISO 27001/27002/27017 Statement of Applicability Certification (27001/27002) Certification (27017)	A.7: Human resource security	March 2023
SOC 1	CA-08: Background checks CA-43: Account revocation	January 3, 2023
SOC 2	CA-07: Standards of Business Conduct (SBC) CA-08: Background checks CA-43: Account revocation ELC-08/13/14: Employment agreements	January 3, 2023

Feedback

Was this page helpful? Yes No

Microsoft pre-employment screening

Article • 03/02/2023

Microsoft works to maintain a safe and secure environment that properly protects our employees, customers, property, personal data, and proprietary information. A comprehensive background screening program is essential for achieving this goal. All candidates for employment and internship are required to satisfactorily complete a background screening prior to commencing employment at Microsoft. Current employees in specific roles may also be subject to periodic rescreening.

The Microsoft Personnel Screening Standard

The Microsoft Personnel Screening Standard applies to all Microsoft employees and interns globally, to the extent lawfully allowable. Personnel involved in the development, operation, or delivery of online services to government or commercial cloud customers may be subject to additional role specific screening requirements and periodic rescreening, in addition to their on-hire screen.

The personnel screening practices for Microsoft online services align with Microsoft's corporate standards and National Institute of Standards and Technology (NIST) controls for personnel screening. Microsoft staff who require access to the following are subject to the Microsoft Personnel Screening Standard:

- Physical access to datacenters, co-locations, secured rooms, cages, server racks, or edge sites that provide the infrastructure supporting online services for government or commercial cloud customers.
- Logical access to government or commercial cloud customer data provided through specific managed environments.
- Network management access to devices and services that transport or store government or commercial cloud customer data.

In countries with robust privacy laws that restrict access to personal information, staff involved in the development, operation, or delivery of online services to government or commercial cloud customers may be subject to additional checks. Specific personnel-related events that trigger screening requirements include:

- New Microsoft staff placed in roles where screening is a defined requirement.
- Internal Microsoft staff transferred or moved to an existing role that currently includes screening as a defined requirement.
- Existing roles that change scope to include screening as a defined requirement.

Microsoft pre-employment screening

Microsoft's Global Security Department is responsible for ensuring candidates are background screened before commencing employment with Microsoft. To the extent permissible by local law the checks include the following:

- SSN trace in the United States or confirmation of identity from a government-issued document in Europe, the Middle East, and Asia.
- Criminal history checks up to seven years for felony and misdemeanor offenses at the state, county, and local levels, and at the federal level if appropriate in the United States. An international criminal history check (dependent on local law) is applicable for areas outside of the United States.
- Confirmation of highest level of academic achievement
- Five-year employment history
- Global sanctions and enforcement check, including Office of Foreign Assets Control (OFAC) list, Bureau of Industry and Security (BIS) list, and Office of Defense Trade Controls Debarred Persons (DDTC) list checks.

Supplier screening

Microsoft requires that Suppliers (contingent staff) conduct pre-placement background checks on all their personnel who will perform services or projects that require any access to Microsoft owned or leased facilities or access to Microsoft resources such as email, network access, card key, or other access badges.

To the extent permissible by local law, Supplier Screening checks include the following items:

- SSN trace in the United States or confirmation of identity from a government-issued document in Europe, the Middle East, and Asia.
- Criminal history checks up to seven years for felony and misdemeanor offenses at the state, county, and local levels, and at the federal level if appropriate in the United States. An international criminal history check (dependent on local law) is applicable for areas outside of the United States.
- Global sanctions and enforcement check

In the United States prior to each placement of each supplier staff member to perform for Microsoft, to the extent permissible by local law, Supplier will conduct a pre-placement, and for certain placements, periodic, background check utilizing Microsoft's preferred background check screening service provider who will confirm to Microsoft that the screening meets the Microsoft standard.

Internationally, the foregoing requirement will be implemented in all countries by September 2020 and where not currently implemented, the contract with the third party specifies Microsoft's requirements for screening that must be conducted and the Supplier is responsible for attesting to Microsoft in writing that the background check has been successfully completed.

Microsoft cloud background check

For candidates hired to provide Microsoft online services in the United States, an additional Microsoft Cloud Background Check is required. For more information, see Microsoft cloud background check.

Microsoft cloud background check

Article • 04/19/2023

A Microsoft Cloud Background Check is required for candidates hired as employees that provide and maintain Microsoft online services in the United States. Microsoft employees in the United States with access to customer data must follow the Microsoft Cloud Background Check process required by all Microsoft online services. Outside of the United States, the process varies. For example, the Microsoft Cloud for Germany uses a Data Trustee approval model, designed to ensure that the Data Trustee (a German company), and not Microsoft, is in control of access to customer data. The Microsoft Cloud in Germany is delivered from datacenters in Germany, and the Operations Centers (OC) containing the technical staff of the Data Trustee are also in Germany. Both the datacenter and the OC facilities are operated, maintained, and controlled by the Data Trustee.

To the extent permissible by local law, the following checks are performed as part of the Microsoft Cloud Background Check.

- United States: Social Security Number trace
- Criminal history checks up to seven years for felony and misdemeanor offenses at the state, county, and local levels, and at the federal level if appropriate in the United States. An international criminal history check (dependent on local law) is applicable for areas outside of the United States.
- Global sanctions and enforcement check, including Office of Foreign Assets Control (OFAC) list, Bureau of Industry and Security (BIS) list, and Office of Defense Trade Controls Debarred Persons (DDTC) list checks.

The results from the Microsoft Cloud Background Check are stored in our employee database and connected to our datacenter access control systems. If the Microsoft Cloud Background Check expires and the employee does not renew it, access to Microsoft services is revoked and no longer available until the Microsoft Cloud Background Check is completed. When the employment relationship with Microsoft ends, all datacenter access is immediately revoked.

All employees with physical or logical access to the environments that run United States Government services are required to have their United States citizenship verified. To verify citizenship, a background screen is ordered by Microsoft for employees and/or new hire candidates in Microsoft's background screening provider's system. The employee and/or new hire candidate will be required to provide an unexpired US Passport which will be validated within the background screening provider's system. The verified results are recorded in Microsoft's HR systems.

Personnel with logical access to Microsoft online services for the U.S. Government Community Cloud are required to comply with federal government requirements of the FBI's Criminal Justice Information Services (CJIS), including personnel screening. CJIS screening in support of the Microsoft 365 U.S. Government service includes a fingerprint-based criminal background check adjudicated by the CJIS system agency designated adjudicator in states that have enrolled in the Microsoft Online Services CJIS support program.

Periodic rescreening

Periodic rescreening and/or additional background checks may be required for certain management, security, or other roles, including employees in roles that require access to customer data.

Resources

- Microsoft Offering: Criminal Justice Information Services (CJIS) Security Policy

Microsoft employee transfer and termination

Article • 03/02/2023

Microsoft, like every other organization, handles employee transfers and terminations as a part of their normal business operation. When an employee changes positions or leaves the company, it is essential to revoke inappropriate access in a timely manner. To facilitate efficient access changes and access revocations, Microsoft uses standardized procedures and automated processes to coordinate the Human Resources Information System (HRIS) with the Identity Management (IDM) system. Automated orchestration between these two systems is essential to maintaining operational consistency, safeguarding Microsoft's online services and data, preventing privilege creep, and reducing risks related to insider threats.

Microsoft online services are designed to operate without standing administrative access to production environments for our engineers. Microsoft uses a Just-In-Time (JIT), Just-Enough-Access (JEA) model to provide engineers with the temporary access needed to support their service when required. To request and use a service team account for JIT access, engineers must request and maintain eligibilities through the IDM tool. When employees are transferred or terminated, their service team account and related eligibilities are automatically modified to prevent inappropriate access.

Transfer and reassignment

Employee transfers are initiated through a transfer transaction request by the employee's manager. The manager creates a requisition and engages with Global Talent Acquisition for the offer letter process. Once the employee accepts the offer for the new role, HR services completes the transfer in the HR core tools, triggering IDM to set an expiration date for all the employee's eligibilities. The employee must submit a request and receive approval from their new manager to retain their eligibilities. Failure to submit a request or receive manager approval results in the revocation of the transferred employee's eligibilities. For transfers that include specific security implications, system accesses and security group memberships are reevaluated immediately to reflect their new role.

Termination

Microsoft uses clearly defined policies and procedures to promptly revoke physical and logical access to Microsoft systems and resources when an employee is terminated. When an employee gives their notice, the employee's manager enters the termination date into the HRIS. Following the employee's last working day, the HRIS marks the employee as terminated and shares the information to IDM, which removes all service team accounts and eligibilities automatically.

For involuntary terminations, HR works with the employee's manager to follow the appropriate steps to terminate and offboard the employee. Similar to a voluntary termination, the termination information is entered into the HRIS along with any necessary steps such as effective date coordination, access removal, and any other steps relative to transitioning out of role.

Privacy & data management overview

Article • 12/06/2023

How does Microsoft approach privacy for customers?

Microsoft's commitment to protecting customer data is set forth in the Product Terms and the Data Protection Addendum (DPA). The foundation of Microsoft's approach to privacy is built on the following principles: customer control, transparency, security, defending data from third party access, no content-based targeting, and compliance with relevant laws and regulations. For more information, see Privacy at Microsoft and the Microsoft Privacy Report to learn more about Microsoft's approach to protecting privacy.

How does Microsoft implement its privacy commitments?

Microsoft maintains the Microsoft Corporate Privacy Policy and Microsoft Privacy Standard to ensure we meet our privacy commitments across the enterprise. Microsoft has governance councils, boards and committees to support the consistent and compliant adoption and implementation of our corporate privacy requirements. Microsoft's privacy program starts with a "hub and spoke" governance model, where responsibility for compliance is shared across the company, some are centralized, and others are distributed. The "hub" of Microsoft's corporate privacy team is the CELA (Corporate, External, and Legal Affairs) group. The "spokes" reside within the company's engineering and functional groups.

Microsoft also has Data Handling Standards in place that provide guidance on how to manage each data classification type within specific activities or scenarios, including requirements to meet the obligations outlined in the Product Terms and DPA.

How does Microsoft collect and process customer data?

The data lifecycle describes how Microsoft processes data based on customer guidance and in compliance with applicable security and privacy laws, as noted in the Product Terms and DPA. Stages of the data lifecycle include collection, processing, storage,

third-party sharing (where applicable), retention, transfer, and deletion. Microsoft's approach to privacy informs each stage of the data lifecycle to protect the privacy of our customers.

Microsoft limits collection of customer data to three data categories: Customer Data, Personal Data, and Professional Services Data as defined in the DPA. Microsoft uses and processes data from these categories to provide products and services in accordance with documented instructions from its customers and for business operations incident to providing the products and services. Specific descriptions of these use and processing activities are defined in the DPA in sections, "Processing to Provide Customer the Products and Services" and "Processing for Business Operations Incident to Providing the Products and Services to Customer," respectively.

How does Microsoft handle third-party sharing?

Third-party sharing is the sharing or onward disclosure of data to third parties. Microsoft will only share data when authorized by the customer or required to do so by applicable law. Microsoft does not give any government (including law enforcement or other government entities) direct or unfettered access to customer data. For more information, see the Law Enforcement Request Report and U.S. National Security Order Report to learn how Microsoft responds to government requests to access data.

Does Microsoft use subprocessors or subcontractors?

For information on how Microsoft manages suppliers, see Supplier Management page.

Who has access to customer data within Microsoft?

To learn how Microsoft manages access to customer data, see Identity and Access Management page.

Where is customer data located?

For Core Online Services, see our commitments outlined in the Product Terms and the DPA to find the most up to date information on location of customer data.

As described in the DPA for the Core Online Services, Microsoft stores Customer Data at rest within certain major geographic areas (each, a Geo) as set forth in the Product Terms. For commercial services in scope for the Microsoft EU Data Boundary, Microsoft stores and processes Customer Data within the European Union as set forth in the Product Terms. Microsoft does not control or limit the regions from which Customer or Customer's end users may access or move Customer Data.

Visit Microsoft Privacy - Where is Your Data Located to learn more.

For Azure and M365 services, visit Azure data residency and M365 data locations.

Other services data locations:

- Azure DevOps Services
- Microsoft Entra ID
- Microsoft Commerce
- Microsoft Defender for Cloud Apps
- Microsoft Defender for Endpoint
- Microsoft Defender for Identity Personal Data Policy
- Microsoft Dynamics 365
- Microsoft Intune
- Microsoft Power Platform
- Microsoft Professional Services
- Microsoft Threat Protection

Microsoft EU Data Boundary

On May 6, 2021, Microsoft announced the EU Data Boundary for the Microsoft Cloud, Microsoft's new commitment to customers in Europe. The EU Data Boundary is a geographically defined boundary within which Microsoft has committed to store and process customer data for our major online services, including Azure, Dynamics 365, Power Platform, and Microsoft 365, subject to limited circumstances where customer data will continue to be transferred outside the EU Data Boundary.

Learn more at:

- Microsoft EU Data Boundary Overview
- What is the EU Data Boundary?

How does Microsoft delete customer data when a customer leaves the service?

The Microsoft Data Handling Standard specifies how long customer data is retained after deletion. When a customer ends their subscription, Microsoft retains customer data in a limited function account for 90 days to enable the customer to extract the data. After the 90-day retention period ends, Microsoft will delete customer data unless authorized to retain it or required to retain it by law. No more than 180 days after expiration or termination of a subscription to Microsoft online services, Microsoft disables the account and deletes all customer data from the account. Once the maximum retention period for any data has elapsed, the data is rendered commercially unrecoverable.

Microsoft also deletes all service-generated and diagnostic data as part of the standard Microsoft data lifecycle unless the data is required to maintain the security and stability of the service. For any subscription, a subscriber can contact Microsoft Support and request expedited subscription de-provisioning. When a customer uses this process, all user data is deleted 3 days after the administrator enters the lockout code provided by Microsoft. This deletion includes data in SharePoint Online and Exchange Online under hold or stored in inactive mailboxes.

Microsoft follows NIST SP-800-88 guidelines for the destruction of devices that are capable of holding data, as described in the Data-bearing device destruction article.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to privacy.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
ISO 27018 Statement of Applicability Certificate	A-2.1: Public cloud PII processor's purpose	April 24, 2023
ISO 27701 Statement of	All controls	April 24, 2023

External audits	Section	Latest report date
Applicability Certificate		
SOC 1	DS-15: Customer subscription termination/expiration SDL-1: Security Development Lifecycle (SDL) methodology LA-4: Protection of confidential customer data	November 17, 2023
SOC 2 SOC 3	DS-15: Customer subscription termination/expiration SDL-1: Security Development Lifecycle (SDL) methodology LA-4: Protection of confidential customer data SOC2-1: Asset classification SOC2-7: Published confidentiality and security obligations	November 17, 2023

Microsoft 365

Expand table

External audits	Section	Latest report date
ISO 27018 Statement of Applicability Certificate	A-2.1: Public cloud PII processor's purpose	March 2023
ISO 27701 Statement of Applicability Certificate	All controls	March 2023
SOC 2	CA-12: Service level agreements (SLAs) CA-17: Microsoft security policy CA-25: Control framework updates	January 3, 2023

Resources

- Microsoft Trust Center: Privacy Principles: Privacy Principles
- Data Residency, Data Sovereignty, and Compliance in the Microsoft Cloud
- Addressing sovereignty requirements with Microsoft Azure

- Protecting Data Privacy Using Microsoft Azure
- Compliance with EU transfer requirements for personal data in the Microsoft Cloud
- Enabling Data Residency and Data Protection in Microsoft Azure Regions
- Microsoft Professional Services privacy and data protection information
- Data Transfer Impact Assessment FAQ

Feedback

Was this page helpful? Yes No

Azure customer data protection

Article • 08/29/2023

Access to customer data by Microsoft operations and support personnel is denied by default. When access to data related to a support case is granted, it is only granted using a just-in-time (JIT) model using policies that are audited and vetted against our compliance and privacy policies. The access-control requirements are established by the following Azure Security Policy:

- No access to customer data, by default.
- No user or administrator accounts on customer virtual machines (VMs).
- Grant the least privilege that's required to complete task; audit and log access requests.

Azure support personnel are assigned unique corporate Active Directory accounts by Microsoft. Azure relies on Microsoft corporate Active Directory, managed by Microsoft Information Technology (MSIT), to control access to key information systems. Multi-factor authentication is required, and access is granted only from secure consoles.

Data protection

Azure provides customers with strong data security, both by default and as customer options.

Data segregation: Azure is a multi-tenant service, which means that multiple customer deployments and VMs are stored on the same physical hardware. Azure uses logical isolation to segregate each customer's data from the data of others. Segregation provides the scale and economic benefits of multi-tenant services while rigorously preventing customers from accessing one another's data.

At-rest data protection: Customers are responsible for ensuring that data stored in Azure is encrypted in accordance with their standards. Azure offers a wide range of encryption capabilities, giving customers the flexibility to choose the solution that best meets their needs. Azure Key Vault helps customers easily maintain control of keys that are used by cloud applications and services to encrypt data. Azure Disk Encryption enables customers to encrypt VMs. Azure Storage Service Encryption makes it possible to encrypt all data placed into a customer's storage account.

In-transit data protection: Microsoft provides a number of options that can be utilized by customers for securing data in transit internally within the Azure network and externally across the Internet to the end user. These include communication through

Virtual Private Networks (utilizing IPsec/IKE encryption), Transport Layer Security (TLS) 1.2 or later (via Azure components such as Application Gateway or Azure Front Door), protocols directly on the Azure virtual machines (such as Windows IPsec or SMB), and more.

Additionally, "encryption by default" using MACsec (an IEEE standard at the data-link layer) is enabled for all Azure traffic traveling between Azure datacenters to ensure confidentiality and integrity of customer data.

Data redundancy: Microsoft helps ensure that data is protected if there is a cyberattack or physical damage to a datacenter. Customers may opt for:

- In-country/region storage for compliance or latency considerations.
- Out-of-country/region storage for security or disaster recovery purposes.

Data can be replicated within a selected geographic area for redundancy but cannot be transmitted outside it. Customers have multiple options for replicating data, including the number of copies and the number and location of replication datacenters.

When you create your storage account, select one of the following replication options:

- **Locally redundant storage (LRS)**: Locally redundant storage maintains three copies of your data. LRS is replicated three times within a single facility in a single region. LRS protects your data from normal hardware failures, but not from a failure of a single facility.
- **Zone-redundant storage (ZRS)**: Zone-redundant storage maintains three copies of your data. ZRS is replicated three times across two to three facilities to provide higher durability than LRS. Replication occurs within a single region or across two regions. ZRS helps ensure that your data is durable within a single region.
- **Geo-redundant storage (GRS)**: Geo-redundant storage is enabled for your storage account by default when you create it. GRS maintains six copies of your data. With GRS, your data is replicated three times within the primary region. Your data is also replicated three times in a secondary region hundreds of miles away from the primary region, providing the highest level of durability. In the event of a failure at the primary region, Azure Storage fails over to the secondary region. GRS helps ensure that your data is durable in two separate regions.

Data destruction: When customers delete data or leave Azure, Microsoft follows strict standards for deleting data, as well as the physical destruction of decommissioned hardware. Microsoft executes a complete deletion of data on customer request and on contract termination. For more information, see Data management at Microsoft.

Customer data ownership

Microsoft does not inspect, approve, or monitor applications that customers deploy to Azure. Moreover, Microsoft does not know what kind of data customers choose to store in Azure. Microsoft does not claim data ownership over the customer information that's entered into Azure.

Records management

Azure has established internal records-retention requirements for back-end data. Customers are responsible for identifying their own record retention requirements. For records that are stored in Azure, customers are responsible for extracting their data and retaining their content outside of Azure for a customer-specified retention period.

Azure allows customers to export data and audit reports from the product. The exports are saved locally to retain the information for a customer-defined retention time period.

Electronic discovery (e-discovery)

Azure customers are responsible for complying with e-discovery requirements in their use of Azure services. If Azure customers must preserve their customer data, they may export and save the data locally. Additionally, customers can request exports of their data from the Azure Customer Support department. In addition to allowing customers to export their data, Azure conducts extensive logging and monitoring internally.

Next steps

To learn more about what Microsoft does to secure the Azure infrastructure, see:

- Azure facilities, premises, and physical security
- Azure infrastructure availability
- Azure information system components and boundaries
- Azure network architecture
- Azure production network
- Azure SQL Database security features
- Azure production operations and management
- Azure infrastructure monitoring
- Azure infrastructure integrity

Exchange Online data deletion in Microsoft 365

Article • 10/20/2023

Within Exchange Online, there are two kinds of deletions: soft deletions and hard deletions. This applies to both mailboxes and items within a mailbox.

Soft-deleted and hard-deleted mailboxes

A soft-deleted user mailbox is a mailbox that has been deleted using the Microsoft 365 admin center or the Remove-Mailbox cmdlet and has still been in the Microsoft Entra ID recycle bin for less than 30 days. A mailbox can become soft-deleted in any of the following ways:

- The user mailbox's associated Microsoft Entra user account is soft-deleted (the user object is out of scope or in the recycle bin container).
- The user mailbox's associated Microsoft Entra user account has been hard-deleted but the Exchange Online mailbox is under a litigation hold or eDiscovery hold.
- The user mailbox's associated Microsoft Entra user account has been purged within the last 30 days; which is the maximum retention length Exchange Online will keep the mailbox in a soft-deleted state before it is permanently purged and unrecoverable.

A hard-deleted user mailbox is a mailbox that has been deleted in one of the following ways:

- The user mailbox has been soft-deleted for more than 30 days, and the associated Microsoft Entra user has been hard-deleted. All mailbox content such as emails, contacts, and files are permanently deleted.
- The user account associated with the user mailbox has been hard-deleted from the Microsoft Entra ID. The user mailbox is now soft-deleted in Exchange Online and stays in a soft-deleted state for 30 days. If in the 30-day period a new Microsoft Entra user is synchronized from the original recipient account with the same **ExchangeGuid** or **ArchiveGuid**, and that new account is licensed for Exchange Online, this will result in a hard deletion of the original user mailbox. All mailbox content such as emails, contacts, and files are permanently deleted.
- A soft-deleted mailbox is deleted using **Remove-Mailbox -PermanentlyDelete**.

The above deletion scenarios assume that the user mailbox isn't in any of the hold states, like Litigation hold or eDiscovery hold. If there is any type of hold on the mailbox,

then the mailbox can't be deleted. For all mail-user recipient types, any Hold settings are ignored and have no effect on hard-deletions or soft-deletions.

Soft-deleted and hard-deleted items

When a user deletes a mailbox item (such as an email message, a contact, a calendar appointment, or a task), the item is moved to the Recoverable Items folder, and into a subfolder named "Deletions". This is referred to as a soft deletion. How long deleted items are kept in the Deletions folder depends on the deleted item retention period that is set for the mailbox. An Exchange Online mailbox keeps deleted items for 14 days by default, but Exchange Online administrators can change this setting to increase the period up to a maximum of 30 days. (For detailed steps to increase the deleted item retention period for an Exchange Online mailbox, see Change how long permanently deleted items are kept for an Exchange Online mailbox.) Users can recover, or purge, deleted items before the retention time for a deleted item expires. To do so, they use the Recover Deleted Items feature in Microsoft Outlook or Outlook on the web.

If a user purges a deleted item by using the Recover Deleted Items feature in Outlook or Outlook on the web, this is known as a hard deletion. In Exchange Online, single item recovery is enabled by default when a new mailbox is created, so an administrator can still recover hard-deleted items before the deleted item retention period expires. Also, if a message is changed by a user or a process, copies of the original item are also retained when single item recovery is enabled.

Page zeroing

Zeroing is a security mechanism that writes either zeros or a binary pattern over deleted data so that the deleted data is more difficult to recover. In Exchange Online, mailbox databases use *pages* as their unit of storage, and implement an overwriting process called *page zeroing*. Page zeroing is enabled by default, and it cannot be disabled by customers or by Microsoft. Page zeroing operations are recorded in the transaction log files so that all copies of a given database are page-zeroed in a similar manner. Zeroing a page on an active database copy causes the page to get zeroed on passive copies of the database.

Page zeroing writes a binary pattern over hard-deleted records. The page-zeroing pattern is specific to Extensible Storage Engine (ESE) operations (the name of the internal database engine used by servers in Exchange Online), and it is different for run-time operations versus background database maintenance operations. (Background database maintenance is a process that continuously checksums and scans each database. Its primary function is to checksum database pages, but it also handles

cleaning up space and zeroing out records and pages that were not zeroed out because of a Store crash.)

The following table lists the fill patterns that correspond to specific run-time operations.

ESE run-time operation	Fill pattern
Replace	R
Record/long value delete	D
Freed page space	H

The following table lists the fill patterns that correspond to specific operations that occur during ESE background database maintenance.

ESE background database maintenance operation	Fill pattern
Record delete	D
Long value delete	L
Freed page space of partially used page	Z
Freed page space of unused page	U

Page zeroing process

The process for page zeroing depends on the deletion scenario. The following table discusses database delete scenarios, and when page zeroing functions occur.

Database delete scenario	ESE process and timeframe to zero database data
Item expires based on the deleted item retention period.	An asynchronous thread writes a binary pattern over the deleted data. This action occurs within milliseconds of the record deletion. If the Store process crashes while the asynchronous zeroing work is still outstanding (or version store cleanup is canceled due to version store growth), the zeroing is completed when background database maintenance processes that section of the database.
View Scenario: Expiration of items from Outlook/Outlook on the web folder view (for example, Conversation view)	Data zeroing occurs when background database maintenance processes that section of the database.
Move Mailbox/Delete Mailbox Scenario: Source mailbox	Data zeroing occurs when background database maintenance processes that section of the database.

Database delete scenario	ESE process and timeframe to zero database data
deleted (expiry of deleted mailbox)	

Mailbox data types without page zeroing

The following mailbox data types have no provisions for page zeroing:

- **Mailbox database transaction logs** - When transaction logs are deleted as part of normal operations, there is no process to zero the blocks in the file system that stored the deleted log file(s). It's likely that the file system will quickly reutilize that free space for newly created logs, but there is no guarantee that this will happen.
- **Content index catalog files** - Exchange Online uses Search Foundation (also known as FAST) for search indexing functionality. The search index catalog is composed of several dozen files stored on the same volume as the mailbox database file. When a message is hard-deleted from the mailbox database, the associated content in the search catalog isn't immediately deleted. Content deletion occurs when Search Foundation does a shadow (or master merge) of many small catalog files in to a single larger file. After the master merge completes, the smaller catalog files are deleted. There is no process to zero the blocks that stored the deleted catalog files.

Continuous replication

Continuous replication (also known as log shipping and replay) is technology in Exchange Online that creates and maintains copies of every mailbox database to provide high availability, site resilience, and disaster recovery. Continuous replication uses the Exchange Server database crash recovery support to provide technology that performs asynchronous updating of one or more copies of a mailbox database. Each mailbox server records database updates made on an active database (for example, user email activity) as log records in a sequential set of 1-MB transaction log files. This set of files is referred to as the log stream. In continuous replication, the log stream is also used to asynchronously update one or more copies of a database. This is accomplished by transmitting the logs to a location containing a passive copy of the active database and then replaying them into the passive database copy. If all logs from the active database are replayed against a passive copy of the database, then the two databases are equivalent, and it is through this process that any physical change made to an active database is replicated to all passive copies of that database.

Any deletion from a mailbox database, whether a mailbox item or an entire mailbox, and whether a soft-delete or a hard-delete, represents a physical change to the active database. Page zeroing also entails making physical changes to the active database. These changes are written to the log files through a process called continuous replication, and when those log files are replayed against passive copies of the database, the same physical changes are made to those passive databases.

SharePoint Online data deletion in Microsoft 365

Article • 05/25/2023

SharePoint Online stores objects as abstracted code within application databases. When a user uploads a file to SharePoint Online, that file is disassembled and translated into application code and stored in multiple tables across multiple databases. In SharePoint Online, all content that a customer uploads is broken into chunks, encrypted (potentially with multiple AES 256-bit keys), and distributed across the datacenter. For specific details about the chunking and encryption process, see Encryption in the Microsoft Cloud.

When using the SharePoint Online application within a browser, items deleted by any user are retained in the site Recycle Bin for 93 days from the time they were deleted them from their original location. They stay in the site Recycle Bin the entire time, unless someone deletes them from there or empties that Recycle Bin. In that case, the items go to the site collection Recycle Bin, also known as the second-stage Recycle Bin, where they stay for the remainder of their retention period. Items deleted from the site collection Recycle Bin are purged immediately. If the item deleted causes the site collection Recycle Bin to exceed its quota, it starts purging the oldest items until there's space for the most recently deleted item. For info about restoring deleted items, see Restore items in the Recycle Bin of a SharePoint site and Restore deleted items from the site collection recycle bin.

When you delete a site collection, you're also deleting the hierarchy of sites in the collection, and all content within them:

- Documents and document libraries
- Lists and list data
- Site configuration settings
- Role and security information that is related to the site or its subsites
- Subsites of the top-level website, their contents, and user information

If you accidentally delete a site collection, it can be restored by a global or SharePoint admin using the SharePoint admin center.

Deleted site collections are retained for 93 days. After 93 days, sites and all their content and settings are permanently deleted, including lists, libraries, pages, and any subsites.

Hard deletion occurs when a user purges deleted items from the site collection Recycle Bin, when the retention and backup periods expire, or when an administrator

permanently deletes a site collection using the Remove-SPODeletedSite cmdlet. When a user hard deletes (permanently deletes, or purges) content from SharePoint Online, all encryption keys for the deleted chunks are also deleted. The blocks on the disks that previously stored the deleted chunks are marked as unused and available for reuse.

Users can also interact with the SharePoint platform by using CSOM/REST APIs. API calls can be used to either recycle or delete items on SharePoint such as files and folders. Recycling an item adds it to the Recycle Bin, similar to the browser application behavior described earlier in this article. However, deleting results in the item being immediately purged and it will not be recoverable from the site or site collection Recycle Bins. You can read more about SharePoint APIs here.

SharePoint Online retains backups of all content for 14 additional days beyond actual deletion to facilitate a full site collection restore for customers. Customers can reach out to Microsoft support to initiate a full site collection or sub-site point in time restore if files have been hard deleted, corrupted, or infected with malware and are unrecoverable using the methods described above. After this 14 day period, data is no longer retained by Microsoft and isn't recoverable.

Data, Privacy, and Security for Microsoft Copilot for Microsoft 365

Article • 12/15/2023

Microsoft Copilot for Microsoft 365 is a sophisticated processing and orchestration engine that provides AI-powered productivity capabilities by coordinating the following components:

- Large language models (LLMs)
- Content in Microsoft Graph, such as emails, chats, and documents that you have permission to access.
- The Microsoft 365 apps that you use every day, such as Word and PowerPoint.

For an overview of how these three components work together, see Microsoft Copilot for Microsoft 365 overview. For links to other content related to Microsoft Copilot for Microsoft 365, see Microsoft Copilot for Microsoft 365 documentation.

> **ⓘ Important**
>
> - Microsoft Copilot for Microsoft 365 is compliant with our existing privacy, security, and compliance commitments to Microsoft 365 commercial customers, including the General Data Protection Regulation (GDPR) and European Union (EU) Data Boundary.
> - Prompts, responses, and data accessed through Microsoft Graph aren't used to train foundation LLMs, including those used by Microsoft Copilot for Microsoft 365.

The information in this article is intended to help provide answers to the following questions:

- How does Microsoft Copilot for Microsoft 365 use your proprietary organizational data?
- How does Microsoft Copilot for Microsoft 365 protect organizational information and data?
- What data is stored about user interactions with Microsoft Copilot for Microsoft 365?
- What data residency commitments does Microsoft Copilot make?
- Can Microsoft Copilot for Microsoft 365 use public web content in its responses?
- What extensibility options are available for Microsoft Copilot for Microsoft 365

- How does Microsoft Copilot for Microsoft 365 meet regulatory compliance requirements?
- Do controls for connected experiences in Microsoft 365 Apps apply to Microsoft Copilot for Microsoft 365?
- Can I trust the content that Microsoft Copilot for Microsoft 365 creates? Who owns that content?
- What are Microsoft's commitments to using AI responsibly?

> **Note**
>
> Microsoft Copilot for Microsoft 365 will continue to evolve over time with new capabilities. To keep up to date on Microsoft Copilot for Microsoft 365 or ask questions, visit the **Microsoft 365 Copilot community** on the Microsoft Tech Community.

How does Microsoft Copilot for Microsoft 365 use your proprietary organizational data?

Microsoft Copilot for Microsoft 365 provides value by connecting LLMs to your organizational data. Microsoft Copilot for Microsoft 365 accesses content and context through Microsoft Graph. It can generate responses anchored in your organizational data, such as user documents, emails, calendar, chats, meetings, and contacts. Microsoft Copilot for Microsoft 365 combines this content with the user's working context, such as the meeting a user is in now, the email exchanges the user had on a topic, or the chat conversations the user had last week. Microsoft Copilot for Microsoft 365 uses this combination of content and context to help provide accurate, relevant, and contextual responses.

> **Important**
>
> Prompts, responses, and data accessed through Microsoft Graph aren't used to train foundation LLMs, including those used by Microsoft Copilot for Microsoft 365.

Microsoft Copilot for Microsoft 365 only surfaces organizational data to which individual users have at least view permissions. It's important that you're using the permission models available in Microsoft 365 services, such as SharePoint, to help ensure the right users or groups have the right access to the right content within your organization. This includes permissions you give to users outside your organization through inter-tenant collaboration solutions, such as shared channels in Microsoft Teams.

When you enter prompts using Microsoft Copilot for Microsoft 365, the information contained within your prompts, the data they retrieve, and the generated responses remain within the Microsoft 365 service boundary, in keeping with our current privacy, security, and compliance commitments. Microsoft Copilot for Microsoft 365 uses Azure OpenAI services for processing, not OpenAI's publicly available services.

> **Note**
>
> When using Microsoft Copilot for Microsoft 365, your organization's data might leave the Microsoft 365 service boundary under the following circumstances:
>
> - When you allow Microsoft Copilot for Microsoft 365 chat experiences to reference public web content. The query sent to Bing might include your organization's data. For more information, see **Microsoft Copilot for Microsoft 365 and public web content**.
> - When you're using plugins to help Microsoft Copilot for Microsoft 365 to provide more relevant information. Check the privacy statement and terms of use of the plugin to determine how it will handle your organization's data. For information, see **Extensibility of Microsoft Copilot for Microsoft 365**.

Abuse monitoring for Microsoft Copilot for Microsoft 365 occurs in real-time, without providing Microsoft any standing access to customer data, either for human or for automated review. While abuse moderation, which includes human review of content, is available in Azure OpenAI, Microsoft Copilot for Microsoft 365 services have opted out of it. Microsoft 365 data isn't collected or stored by Azure OpenAI.

> **Note**
>
> We may use customer feedback, which is optional, to improve Microsoft Copilot for Microsoft 365, just like we use customer feedback to improve other Microsoft 365 services and Microsoft 365 apps. We don't use this feedback to train the foundation LLMs used by Microsoft Copilot for Microsoft 365. Customers can manage feedback through admin controls. For more information, see **Manage Microsoft feedback for your organization** and **Providing feedback about Microsoft Copilot for Microsoft 365**.

Data stored about user interactions with Microsoft Copilot for Microsoft 365

When a user interacts with Microsoft Copilot for Microsoft 365 apps (such as Word, PowerPoint, Excel, OneNote, Loop, or Whiteboard), we store data about these interactions. The stored data includes the user's prompt, how Copilot responded, and information used to ground Copilot's response. For example, this stored data provides users with Copilot interaction history in Microsoft 365 Chat and meetings in Microsoft Teams. This data is processed and stored in alignment with contractual commitments with your organization's other content in Microsoft 365. The data is encrypted while it's stored and isn't used to train foundation LLMs, including those used by Microsoft Copilot for Microsoft 365.

To view and manage this stored data, admins can use Content search or Microsoft Purview. Admins can also use Microsoft Purview to set retention policies for the data related to chat interactions with Copilot. For more information, see the following articles:

- Overview of Content search
- Microsoft Purview data security and compliance protections for Microsoft Copilot
- Learn about retention for Copilot for Microsoft 365

For Microsoft Teams chats with Copilot, admins can also use Microsoft Teams Export APIs to view the stored data.

Deleting the history of user interactions with Microsoft Copilot for Microsoft 365

To delete a user's history of interactions with Microsoft Copilot for Microsoft 365, which includes user prompts and the responses Copilot returns, Microsoft 365 admins can submit an online support ticket in the Microsoft 365 admin center. In this ticket, admins should include their Tenant ID and the user's Object ID for which they want data deleted. The ticket will mark the history for permanent, hard-deletion. For any new requests, open a new ticket with support and avoid editing your existing requests.

Microsoft Copilot for Microsoft 365 and the EU Data Boundary

Microsoft Copilot for Microsoft 365 calls to the LLM are routed to the closest data centers in the region, but also can call into other regions where capacity is available during high utilization periods.

For European Union (EU) users, we have additional safeguards to comply with the EU Data Boundary. EU traffic stays within the EU Data Boundary while worldwide traffic can be sent to the EU and other countries or regions for LLM processing.

Microsoft Copilot for Microsoft 365 and data residency

Customers with Advanced Data Residency (ADR) in Microsoft 365 or Microsoft 365 Multi-Geo can purchase and enable Microsoft Copilot for Microsoft 365. At this time, Microsoft doesn't provide data residency commitments for Microsoft Copilot for Microsoft 365, beyond EU Data Boundary. When customers store data generated by Copilot in Microsoft 365 products that have data residency commitments under the Product Terms, the applicable commitments will be upheld.

Microsoft Copilot for Microsoft 365 and public web content

Microsoft Copilot Graph-grounding chat experiences can reference public web content from the Bing search index to ground user prompts and responses. Based on the user's prompt, Copilot for Microsoft 365 determines whether it needs to use Bing to query public web content to help provide a relevant response to the user. There are controls available to manage the use of public web content for both admins and users.

> **Note**
>
> Public web content grounding in Copilot uses only the Bing Search service. Copilot with commercial data protection (previously named Bing Chat Enterprise) is a separate offering and not involved with public web content grounding.

Details on how public web content grounding works

When public web grounding is enabled, Copilot for Microsoft 365 may automatically generate a web search query, if Copilot for Microsoft 365 determines that web data can improve the quality of the response. The search query is based on the user's prompt, Copilot interaction history, and relevant data the user has access to in Microsoft 365. This web search query might be displayed to the user after the prompt is entered. For example, the user might see the phrase "searching for..." followed by the search query. The query is passed to the Bing Search API, which is part of the Bing Search service, to retrieve information from the web to ground a response.

Web search queries might not contain all the words from a user's prompt. They're generally based off a few terms used to find relevant information on the web. However, they may still include some confidential data, depending on what the user included in

the prompt. Queries sent to the Bing Search API by Copilot for Microsoft 365 are disassociated from the user ID or tenant ID.

Once web data is received, Copilot for Microsoft 365 passes the web data to the LLM to generate a richer response by including the latest information from the web and any relevant citations. In this process, the user's prompts and Copilot's responses remain within the Microsoft 365 service boundary. Only the search query, which is abstracted from the user's prompt and grounding data, goes to the Bing Search API outside the boundary.

Microsoft Bing is a separate business from Microsoft 365 and data is managed independently of Microsoft 365. The use of Bing is covered by the Microsoft Services Agreement between each user and Microsoft, together with the Microsoft Privacy Statement. The Microsoft Products and Services Data Protection Addendum (DPA) doesn't apply to the use of Bing.

Controls available to manage the use of public web content

Admins can prevent their users from referencing public web content in their requests. For more information, see Manage access to public web content in Microsoft Copilot for Microsoft 365 responses. Even when allowed by the admin, users still have the option whether or not they want to reference public web content in their requests. For more information, see Use additional data sources with Microsoft 365 Copilot.

> **Note**
>
> The policy settings that control the use of optional connected experiences in Microsoft 365 Apps don't apply to Microsoft Copilot for Microsoft 365 and public web content.

Extensibility of Microsoft Copilot for Microsoft 365

While Microsoft Copilot for Microsoft 365 is already able to use the apps and data within the Microsoft 365 ecosystem, many organizations still depend on various external tools and services for work management and collaboration. Microsoft Copilot for Microsoft 365 experiences can reference third-party tools and services when responding to a user's request by using Microsoft Graph connectors or plugins. Data from Graph

connectors can be returned in Microsoft Copilot for Microsoft 365 responses if the user has permission to access that information.

When plugins are enabled, Microsoft Copilot for Microsoft 365 determines whether it needs to use a specific plugin to help provide a relevant response to the user. If a plugin is needed, Microsoft Copilot for Microsoft 365 generates a search query to send to the plugin on the user's behalf. The query is based on the user's prompt, Copilot interaction history, and data the user has access to in Microsoft 365.

In the **Integrated apps** section of the Microsoft 365 admin center, admins can view the permissions and data access required by a plugin as well as the plugin's terms of use and privacy statement. Admins have full control to select which plugins are allowed in their organization. A user can only access the plugins that their admin allows and that the user installed or is assigned. Microsoft Copilot for Microsoft 365 only uses plugins that are turned on by the user.

> **Note**
>
> The policy settings that control the use of optional connected experiences in Microsoft 365 Apps don't apply to plugins.

For more information, see the following articles:

- Manage Plugins for Copilot in Integrated Apps
- Extend Microsoft Copilot for Microsoft 365
- How Microsoft Copilot for Microsoft 365 can work with your external data

How does Microsoft Copilot for Microsoft 365 protect organizational data?

The permissions model within your Microsoft 365 tenant can help ensure that data won't unintentionally leak between users, groups, and tenants. Microsoft Copilot for Microsoft 365 presents only data that each individual can access using the same underlying controls for data access used in other Microsoft 365 services. Semantic Index honors the user identity-based access boundary so that the grounding process only accesses content that the current user is authorized to access. For more information, see Microsoft's privacy policy and service documentation.

When you have data that's encrypted by Microsoft Purview Information Protection, Microsoft Copilot for Microsoft 365 honors the usage rights granted to the user. This encryption can be applied by sensitivity labels or by restricted permissions in Microsoft

365 apps by using Information Rights Management (IRM). For more information about using Microsoft Purview with Microsoft Copilot for Microsoft 365, see Microsoft Purview data security and compliance protections for Microsoft Copilot.

We already implement multiple forms of protection to help prevent customers from compromising Microsoft 365 services and applications or gaining unauthorized access to other tenants or the Microsoft 365 system itself. Here are some examples of those forms of protection:

- Logical isolation of customer content within each tenant for Microsoft 365 services is achieved through Microsoft Entra authorization and role-based access control. For more information, see Microsoft 365 isolation controls.
- Microsoft uses rigorous physical security, background screening, and a multi-layered encryption strategy to protect the confidentiality and integrity of customer content.
- Microsoft 365 uses service-side technologies that encrypt customer content at rest and in transit, including BitLocker, per-file encryption, Transport Layer Security (TLS) and Internet Protocol Security (IPsec). For specific details about encryption in Microsoft 365, see Encryption in the Microsoft Cloud.
- Your control over your data is reinforced by Microsoft's commitment to comply with broadly applicable privacy laws, such as the GDPR, and privacy standards, such as ISO/IEC 27018, the world's first international code of practice for cloud privacy.
- For content accessed through Microsoft Copilot for Microsoft 365 plug-ins, encryption can exclude programmatic access, thus limiting the plug-in from accessing the content. For more information, see Configure usage rights for Azure Information Protection.

Meeting regulatory compliance requirements

As regulation in the AI space evolves, Microsoft will continue to adapt and respond to fulfill future regulatory requirements.

Microsoft Copilot for Microsoft 365 is built on top of Microsoft's current commitments to data security and privacy in the enterprise. There's no change to these commitments. Microsoft Copilot for Microsoft 365 is integrated into Microsoft 365 and adheres to all existing privacy, security, and compliance commitments to Microsoft 365 commercial customers. For more information, see Microsoft Compliance.

Beyond adhering to regulations, we prioritize an open dialogue with our customers, partners, and regulatory authorities to better understand and address concerns, thereby fostering an environment of trust and cooperation. We acknowledge that privacy, security, and transparency aren't just features, but prerequisites in the AI-driven landscape at Microsoft.

Additional information

Microsoft Copilot for Microsoft 365 and policy settings for connected experiences

If you turn off connected experiences that analyze content for Microsoft 365 Apps on Windows or Mac devices in your organization, Microsoft Copilot for Microsoft 365 features won't be available to your users in the following apps:

- Excel
- PowerPoint
- OneNote
- Word

Similarly, Microsoft Copilot for Microsoft 365 features in those apps on Windows or Mac devices won't be available if you turn off the use of connected experiences for Microsoft 365 Apps.

For more information about these policy settings, see the following articles:

- Use policy settings to manage privacy controls for Microsoft 365 Apps for enterprise *(for Windows)*
- Use preferences to manage privacy controls for Office for Mac

About the content that Microsoft Copilot for Microsoft 365 creates

The responses that generative AI produces aren't guaranteed to be 100% factual. While we continue to improve responses, users should still use their judgment when reviewing the output before sending them to others. Our Microsoft Copilot for Microsoft 365 capabilities provide useful drafts and summaries to help you achieve more while giving you a chance to review the generated AI rather than fully automating these tasks.

We continue to improve algorithms to proactively address issues, such as misinformation and disinformation, content blocking, data safety, and preventing the

promotion of harmful or discriminatory content in line with our responsible AI principles.

Microsoft doesn't claim ownership of the output of the service. That said, we don't make a determination on whether a customer's output is copyright protected or enforceable against other users. This is because generative AI systems may produce similar responses to similar prompts or queries from multiple customers. Consequently, multiple customers may have or claim rights in content that is the same or substantially similar.

If a third party sues a commercial customer for copyright infringement for using Microsoft's Copilots or the output they generate, we'll defend the customer and pay the amount of any adverse judgments or settlements that result from the lawsuit, as long as the customer used the guardrails and content filters we have built into our products. For more information, see Microsoft announces new Copilot Copyright Commitment for customers.

Committed to responsible AI

As AI is poised to transform our lives, we must collectively define new rules, norms, and practices for the use and impact of this technology. Microsoft has been on a Responsible AI journey since 2017, when we defined our principles and approach to ensuring this technology is used in a way that is driven by ethical principles that put people first.

At Microsoft, we're guided by our AI principles, our Responsible AI Standard, and decades of research on AI, grounding, and privacy-preserving machine learning. A multidisciplinary team of researchers, engineers and policy experts reviews our AI systems for potential harms and mitigations — refining training data, filtering to limit harmful content, query- and result-blocking sensitive topics, and applying Microsoft technologies like InterpretML and Fairlearn to help detect and correct data bias. We make it clear how the system makes decisions by noting limitations, linking to sources, and prompting users to review, fact-check, and adjust content based on subject-matter expertise. For more information, see Governing AI: A Blueprint for the Future.

We aim to help our customers use our AI products responsibly, sharing our learnings, and building trust-based partnerships. For these new services, we want to provide our customers with information about the intended uses, capabilities, and limitations of our AI platform service, so they have the knowledge necessary to make responsible deployment choices. We also share resources and templates with developers inside organizations and with independent software vendors (ISVs), to help them build effective, safe, and transparent AI solutions.

Related articles

- Microsoft Copilot for Microsoft 365 requirements
- Get started with Microsoft Copilot for Microsoft 365
- Microsoft Copilot adoption site

Feedback

Was this page helpful? Yes No

Data immutability in Microsoft 365

Article • 04/19/2023

Regulatory compliance, internal governance requirements, or litigation risks require organizations to preserve email and associated data in a discoverable form. All data in the system must be discoverable and none of it can be destroyed or altered. The industry-standard term for this is "immutability."

Traditional methods for immutability typically worked by moving email messages to a separate, read-only storage location. While such systems serve the purpose of preserving mailbox items for discovery, they often affect the user experience by removing preserved items from the customary daily workflow. For IT professionals, this immutability approach requires the deployment and ongoing maintenance of a separate server and storage infrastructure. Discovery is performed with tools external to the mail system and includes associated deployment and maintenance costs.

Through in-place retention and preservation policy features of archiving in Microsoft 365 and its services, you can preserve and retain many classes of incoming, internal, and outgoing data. This includes:

- Inbound and outbound email communications
- Books and records contained in email form or in shared online documents
- Meeting requests
- Faxes
- Instant messages
- Documents shared during online meetings
- Voicemails

In addition, Microsoft has developed add-on features to allow archiving of data from other sources through integration with third-party data capturing and management solutions. After third-party data is imported, you can apply Microsoft Purview features to the data, including:

- Litigation Hold
- In-Place eDiscovery and Hold
- Compliance Search
- In-Place Archiving
- Mailbox auditing
- Retention Policies

For example, when a mailbox is placed on Litigation Hold, third-party data is preserved. You can search third-party data using In-Place eDiscovery or Compliance Search. Or you

can apply archiving and retention policies to third-party data just like you can for Microsoft data. Archiving third-party data in Microsoft 365 helps your organization stay compliant with government and regulatory policies.

Archiving in Microsoft 365 provides Securities and Exchange Commission (SEC) Rule 17a-4-compliant storage. Microsoft 365 preserves permanent files of all data collected in a non-rewriteable, non-erasable format using in-place retention policies and preservation policies, including preservation lock.

Specifically:

- All records stored using the retention policies noted above are retained in a dedicated storage area out of the purview of the ordinary user. Only authorized users can access and search these records, but can't alter or erase them.
- Metadata for each item includes a timestamp used in the calculation of retention duration. Timestamps are applied when a new item is received or created and can't be modified or removed from the metadata.
- Archiving in Microsoft 365 allows users to combine different retention policies and hold actions to create granular retention policies. These policies define the type or location of the items preserved and the duration of preservation.
- The Preservation Lock feature allows users to choose whether to make the policy a restrictive policy. A restrictive policy prohibits anyone from having the ability to remove, disable, or make any changes to the retention policy. This means that once Preservation Lock is enabled, it can't be disabled, and no mechanism will exist under which any data from existing custodians that has been collected by the retention policies in place may be overwritten, modified, erased, or deleted during the preservation period. Further, the hold period set by Preservation Lock may not be shortened or decreased. It may, however, be lengthened, if a legal requirement to continue retention of the stored data, as noted above. Preservation Lock ensures that no one, not even administrators or those with certain control access, may change the settings or overwrite or erase data that has been stored, bringing archiving in Microsoft 365 in line with the guidance set forth in the 2003 Release of SEC Rule 17a-4.

To understand how Microsoft 365 helps you meet regulatory obligations, specifically in relation to Rule 17a-4 requirements, see the Regulatory requirements for information governance and records management article and linked Cohasset Assessment. The Cohasset Assessment also provides an in-depth analysis of Microsoft 365 archiving features and functionalities against each of the requirements under SEC Rule 17a-4 and demonstrates to regulated customers how Microsoft 365 archiving can enable them to meet these requirements.

Where your Microsoft 365 customer data is stored

Article • 10/31/2023

> **Note**
>
> The **Israel** local data center region launched on November 1, 2023. If your organization requires the migration of your Microsoft 365 customer data to Israel, and data residency commitments for Israel, see **Advanced Data Residency**.

> **Note**
>
> The **Italy** local data center region launched on October 3, 2023. If your organization requires the migration of your Microsoft 365 customer data to Italy, and data residency commitments for Italy, see **Advanced Data Residency**.

> **Note**
>
> For tenants in Australia, Brazil, Canada, France, Germany, India, Japan, Qatar, South Korea, Norway, South Africa, Sweden, Switzerland, United Arab Emirates, and the United Kingdom, additional workloads are available for data residency commitments. For more information, see **Advanced Data Residency**.

See the following links to understand how you can determine current data residency and data residency commitments.

- Exchange Online Data Location
- SharePoint Online (ODSP) and OneDrive for Business Data Location
- Microsoft Teams Data Location
- Microsoft Defender for Office (MDO P1) Data Location
- Office for the Web (Office Online) Data Location
- Viva Connections Data Location
- Viva Topics Data Location

- Microsoft Purview (select services) Data Location
 - Audit (Standard)
 - Audit (Premium)
 - Data loss prevention
 - Data retention
 - Information barriers
 - Office message encryption
 - Records management
 - Sensitivity labels
- Microsoft Entra ID Data Location
- Whiteboard Data Location
- Forms Data Location
- Intune Data Location
- Planner Data Location
- Viva Goals Data Location
- Viva Insights – Advanced, Mgr, Leader Data Location
- Viva Insights – Personal Data Location
- Viva Learning Data Location
- Viva Pulse Data Location
- Yammer Data Location
- Viva Engage Data Location
- Office for mobile Data Location
- OneNote Services Data Location
- Power Apps for Microsoft 365 Data Location
- Stream Data Location
- Shifts Data Location

Data retention, deletion, and destruction in Microsoft 365

Article • 08/07/2023

Microsoft has a Data Handling Standard policy for Microsoft 365 that specifies how long customer data is retained after deletion. There are generally two scenarios in which customer data is deleted:

- **Active Deletion**: The tenant has an active subscription and a user or administrator deletes data, or administrators delete a user.
- **Passive Deletion**: The tenant subscription ends.

Data Retention

For each of these deletion scenarios, the following table shows the maximum data retention period, by data category and classification:

Data Category	Data Classification	Description	Examples	Retention Period
Customer Data	Customer Content	Content directly provided/created by admins and users Includes all text, sound, video, image files, and software created and stored in Microsoft data centers when using the services in Microsoft 365	Examples of the most commonly used Microsoft 365 applications that allow users to author data include Word, Excel, PowerPoint, Outlook, and OneNote Customer content also includes customer-owned/provided secrets (passwords, certificates, encryption keys, storage keys)	**Active Deletion Scenario:** at most 30 days **Passive Deletion Scenario:** at most 180 days
Customer Data	End User Identifiable Information (EUII)	Data that identifies or could be used to identify the user of a Microsoft service. EUII does not contain Customer content	User name or display name (DOMAIN\UserName) User principal name (name@domain)	**Active Deletion Scenario:** at most 180 days (only a tenant administrator action)

Data Category	Data Classification	Description	Examples	Retention Period
			User-specific IP addresses	**Passive Deletion Scenario:** at most 180 days
Personal Data (data not included in Customer Data)	End User Pseudonymous Identifiers (EUPI)	An identifier created by Microsoft tied to the user of a Microsoft service. When combined with other information, such as a mapping table, EUPI identifies the end user EUPI does not contain information uploaded or created by the customer	User GUIDs, PUIDs, or SIDs Session IDs	**Active Deletion Scenario:** at most 30 days **Passive Deletion Scenario:** at most 180 days

Subscription Retention

In the term of an active subscription, a subscriber can access, extract, or delete customer data stored in Microsoft 365. If a paid subscription ends or is terminated, Microsoft retains customer data stored in Microsoft 365 in a limited-function account for 90 days to enable the subscriber to extract the data. After the 90-day retention period ends, Microsoft disables the account and deletes the customer data. No more than 180 days after expiration or termination of a subscription to Microsoft 365, Microsoft disables the account and deletes all customer data from the account. Once the maximum retention period for any data has elapsed, the data is rendered commercially unrecoverable.

For a free trial, your account moves into a grace status for 30 days in most countries and regions. During this grace period, you can purchase Microsoft 365. If you decide not to buy Microsoft 365, you can either cancel your trial or let the grace period expire, and your trial account information and data is deleted.

Related Links

- Data-bearing device destruction
- Immutability in Office 365

- Exchange Online Data Deletion
- SharePoint Online Data Deletion
- Microsoft Defender for Endpoint and Defender for Business data storage and privacy

Create a well-designed data classification framework

Article • 03/02/2023

As you develop, revamp, or refine your data classification framework, consider the following leading practices:

- **Do not expect to go from 0-100 on day 1**: Microsoft recommends a crawl-walk-run approach, prioritizing features critical to the organization and mapping them against a timeline. Complete the first step, ensure it was successful, and then move on to the next phase applying lessons learned. Remember that your organization may still be exposed to risk while you design your data classification framework, so it is okay to start small with just a few classification levels and expand later as needed.

- **You are not just writing for cybersecurity professionals**: Data classification frameworks are meant for a broad audience, including your average staff member, your legal and compliance teams, and your IT team. It is important to write clear, easy-to-understand definitions for your data classification levels, providing real-world examples wherever possible. Try to avoid jargon and consider a glossary for acronyms and highly technical terms. For example, use 'Personally Identifiable Information' and provide a definition instead of simply saying 'PII'.

- **Data classification frameworks are meant to be implemented**: For data classification frameworks to be successful, they must be implemented. It is especially relevant when crafting the control requirements for each data classification level. Make sure requirements are clearly defined and that they anticipate and address any ambiguity that might arise during implementation. For example, if you have a control around Personally Identifiable Information, make sure to spell out exactly what that control means, such as Social Security or passport number.

- **Only go granular if you need to**: Data classification frameworks typically contain anywhere from 3-5 data classification levels. But just because you *can* include five levels does not mean you *should*. Consider the following criteria when deciding on the number of classification levels you need:
 - Your industry and your associated regulatory obligations (highly regulated industries tend to need more classification levels)
 - The operational overhead required to maintain a more complex framework

 - Your users and their ability to comply with the increased complexity and nuance associated with more classification levels
 - User experience and accessibility when seeking to apply manual classification across multiple device types

- **Get the right people involved**: Having a senior stakeholder is critical for success, as many projects struggle to start or take longer without senior management backing. Data classification frameworks are typically owned by information technology teams, but they may have legal, compliance, privacy, and change management implications. To ensure you are creating a framework that helps protect your business, be sure to include privacy and legal stakeholders such as your Chief Privacy Officer and the Office of General Counsel in the development of your policy. If your organization has a Compliance division, information governance professionals, or a records management team, they may also have valuable input. As your framework is rolled out to the business, your Communications department also has a key role to play for internal messaging and adoption.

- **Balance security against convenience**: A common mistake is to draft a secure but overly restrictive data classification framework. This framework may have been designed with security in mind but is often difficult to implement in practice. If users need to follow complex, rigid, and time-consuming procedures to apply the framework in their daily lives, there always is a risk that they may no longer believe in its value and will eventually stop following procedures. This risk exists at all levels of the organization, including executive-level (C-suite) managers within the organization. A good balance of security against convenience alongside easy-to-use tools usually lead to wider user adoption and use. If there are gaps in your framework, do not wait until everything is perfect to start implementation. Instead, assess the risk or gap, create a plan to mitigate, and continue moving forward. Remember that information protection is a journey, it is not something that is activated overnight and then done. Plan, implement some capabilities, confirm success, and iterate to the next milestone as tools evolve and users gain maturity and experience.

Also keep in mind that a data classification framework only addresses *what* your organization should do to protect sensitive data. Data classification frameworks are often accompanied by data handling rules or guidelines that define *how* to put these policies in place from a technical and technology perspective. In the following sections, we turn to some practical guidance on how to take your data classification framework from a policy document to a fully implemented and actionable initiative.

Pain points in creating a data classification framework

Data classification efforts are by nature wide-reaching, touching nearly every business function within an enterprise. Because of this broad scope and the complexity of managing content in modern digital environments, companies often face challenges in knowing where to start, how to manage a successful implementation, and how to measure their progress. Common pain points often include:

- Designing a robust and easy-to-understand data classification framework, including determining classification levels and associated security controls.
- Developing an implementation plan that includes confirming the appropriate technology solution, aligning the plan to existing business processes, and identifying impact to the workforce.
- Setting up a data classification framework within the chosen technology solution and addressing any gaps between the technology capabilities of the tool and the framework itself.
- Establishing a governance structure that oversees the on-going maintenance and health of data classification efforts.
- Identifying specific key performance indicators (KPIs) to monitor and measure progress.
- Increasing awareness and understanding of data classification policies, why they are important, and how to comply with them.
- Complying with internal audit reviews that target data loss and cybersecurity controls.
- Training and engaging users so that they become mindful of the need for correct classification in their daily work and apply the right classification measures.

Change management and training

Organizations today use tools such as Microsoft 365 to implement their data classification framework. The purpose is to try to automate the classification of data and not increase the burden on your workforce. This structure does not mean that your organization has no responsibility to increase awareness of the need to manage content and protect the organization from the risks discussed in this paper. The leading practice continues to be to conduct awareness training across the organization as part of the annual training schedule. Our experience shows that putting robust and comprehensive effort into training your users, who are the key audience performing this work, increases their 'buy-in' to the effort and can increase adoption and quality. Adding label recommendations and in-app tips can amplify these efforts. This training does not need

to be an extensive standalone course. Your organization may incorporate it into other regular training such as your information security annual training and then include an overview of data classification levels and definitions. The main point is that your workforce has the understanding that even though the tool is automating the classifying of data, that does not eliminate each user's overall responsibility for protecting the data in accordance with your company policy.

In addition, you should consider more in-depth training for IT and information security teams to reinforce operational readiness. The teams that manage the tool and the data classification framework must be on the same page. This coordination may require you to invest in a more robust training schedule that may be more often than annually. Investment in more frequent training represents another avenue to reduce risk to your organization. This team is responsible for the implementation and therefore could be a point of failure if not trained on the tool and the policy.

If you need to manually tag content in the tool, developing a group of super-users that have received more advanced training is appropriate. These super users would be engaged for situations where users are required to manually tag documents with data sensitivity labels and would have deep understanding of your organization's data classification framework and regulatory requirements.

Finally, your leadership should prioritize the championing of information security behaviors to reinforce to the workforce the importance of risk management initiatives. These include developing and implementing a robust data classification framework and assigning key leaders to promote the initiative, sometimes referred to as ambassadors or champions of the change.

Governance and maintenance

After you have developed and implemented your data classification framework, ongoing governance and maintenance will be critical to your success. In addition to tracking how sensitivity labels are used in practice, you will need to update your control requirements based on changes in regulations, cybersecurity leading practices, and the nature of the content you manage. Governance and maintenance efforts may include:

- Establishing a governance body dedicated to data classification or adding a data classification responsibility to the charter of an existing information security body.
- Defining roles and responsibilities for users overseeing data classification.
- Establishing KPIs to monitor and measure progress.
- Tracking cybersecurity leading practices and regulatory changes.
- Developing standard operating procedures that support and enforce a data classification framework.

Industry considerations

While the basic principles for developing a strong data classification framework are universal, the details of your framework will depend on the nature of your industry and the unique compliance and security factors your data demands.

For example, financial services firms may need to consider compliance with several regulatory frameworks depending on the scope of their business and the regions in which they operate. Securities firms in the United States must conform to account regulations like SEC Rule 17a-4(f) or FINRA Rule 4511 which address requirements around the security and retention of books and records. Similarly, firms operating in the United Kingdom need to consider FCA compliance.

Government agencies face various regulations governing their data, which vary based on territory and the nature of their work. In the United States, for instance, government agencies and their agents that access federal tax information (FTI) are subject to IRS 1075, which aims to minimize the risk of loss, breach, or misuse of federal tax information.

While financial services firms and government agencies are among the most heavily regulated organizations in the world, most businesses have industry-specific considerations that need to be considered. Some examples include:

- Health industry organizations ensuring compliance with HIPAA.
- Education institutions, from K-12 schools to universities, managing FERPA compliance.
- Drug manufacturers working to comply with GxP guidelines in their country or region around information security.
- Media, retail, and many other companies dealing with GDPR compliance.
- Delivery and storage of entertainment, software, and information content dealing with CDSA.
- Energy industry information security complying with NERC CIP standard.

Implementing your data classification framework in Microsoft 365

Once you have developed your data classification framework, your next step is implementation. The Microsoft Purview compliance portal allows administrators to discover, classify, review, and monitor their data in accordance with their data classification framework. Sensitivity labels can be used to help protect your data by enforcing various protections such as encryption and content marking. They can be

applied to data manually; by default, based on policy settings; or automatically, as the result of a condition such as identified PII.

For smaller organizations or organizations with a relatively streamlined data classification framework, creating a single sensitivity label for each of your data classification levels may suffice. The following example shows a one-to-one data classification level to sensitivity label mapping:

Classification label	Sensitivity label	Label settings	Published to
Unrestricted	Unrestricted	Apply 'Unrestricted' footer	All users
General	General	Apply 'General' footer	All users

> **Tip**
>
> During a Microsoft internal information protection pilot, there were difficulties with the understanding and use of the 'Personal' label. Users were confused as to whether this meant PII or merely related to a personal matter. The label was changed to 'non-business' to be clearer. This example shows that taxonomy does not need to be perfect from the start. Start with what you think is right, pilot it, and adjust the label based on feedback if needed

For larger organizations with a global reach or more complex information security needs, you may find this one-to-one relationship between the number of classification levels in your policy and the number of sensitivity labels in your Microsoft 365 environment to be a challenge. This challenge is especially true in global organizations where a given data classification level such as 'Restricted' may have a different definition or different set of controls depending on region.

For more information on implementation, see Understand data classification and Learn about sensitivity labels.

References

- Microsoft compliance offerings

Data classification & sensitivity label taxonomy

Article • 03/02/2023

Sensitive data presents significant risk to a company if it is stolen, inadvertently shared, or exposed through a breach. Risk factors include reputational damage, financial impact, and loss of competitive advantage. Protecting the data and information your business manages is a top priority for your organization, but you may find it difficult to know if your efforts are truly effective, given the amount of content held by your enterprise.

In addition to volume, your content may range in importance from highly sensitive and impactful to trivial and transient. It can also be under the purview of various regulatory compliance requirements. Knowing what to prioritize and where to apply controls can be a challenge. Read on to learn about *data classification*, an important tool for protecting your content from theft, sabotage, or inadvertent destruction, and how Microsoft 365 can help you meet your information security goals.

What is data classification?

Data classification is a specialized term used in the fields of cybersecurity and information governance to describe the process of identifying, categorizing, and protecting content according to its sensitivity or impact level. In its most basic form, data classification is a means of protecting your data from unauthorized disclosure, alteration, or destruction based on how sensitive or impactful it is.

What is a data classification framework?

Often codified in a formal, enterprise-wide policy, a data classification framework (sometimes called a 'data classification policy') is typically comprised of 3-5 classification levels. These usually include three elements: a name, description, and real-world examples. Microsoft recommends no more than five top-level parent labels, each with five sub-labels (25 total) to keep the user interface (UI) manageable. Levels are typically arranged from least to most sensitive such as *Public*, *Internal*, *Confidential*, and *Highly Confidential*. Other level name variations you may encounter include *Restricted*, *Unrestricted*, and *Consumer Protected*. Microsoft recommends label names that are self-descriptive and that highlight their relative sensitivity clearly. For example, *Confidential* and *Restricted* may leave users guessing which label is appropriate, while *Confidential* and *Highly Confidential* are clearer on which is more sensitive.

The following table shows an example of a *Highly Confidential* data classification framework level:

Classification level	Description	Examples
Highly Confidential	Highly Confidential data is the most sensitive type of data stored or managed by the enterprise and may require legal notifications if breached or otherwise disclosed. Restricted Data requires the highest level of control and security, and access should be limited to "need-to- know."	Sensitive Personally Identifiable Information (Sensitive PII) Cardholder Data Protected Health Information (PHI) Bank Account Data

> **Tip**
>
> Microsoft's corporate data classification framework originally used a category and label named 'Internal' during pilot phase but found that there were legitimate reasons for a document to be shared externally and shifted to using 'General'.

Another important component of a data classification framework is the controls associated with each level. Data classification levels by themselves are simply labels (or tags) that indicate the value or sensitivity of the content. To *protect* that content, data classification frameworks define the controls that should be in place for each of your data classification levels. These controls may include requirements related to:

- Storage type and location
- Encryption
- Access control
- Data destruction
- Data loss prevention
- Public disclosure
- Logging and tracking access
- Other control objectives, as needed

Your security controls will vary by data classification level, such that the protective measures defined in your framework increase commensurate with the sensitivity of your content. For example, your data storage control requirements will vary depending upon the media that is being used as well as upon the classification level applied to a given piece of content. The following table shows an example of data classification controls for a specific storage type:

Storage type	Confidential	Internal	Unrestricted
Removable Storage	Prohibited	Prohibited unless encrypted	No control required

Correctly applying the right level of data classification can be complex in real-life situations and may sometimes overwhelm end users. Once a policy or standard has been created that defines the required levels of data classification, it is important to guide end users on how to bring this framework to life in their daily work. This area is where data classification handling rules or guidelines come in.

Data classification handling guidelines will help end users with specific guidance on how to handle each level of data appropriately, for different storage media throughout their lifecycle. These guidelines help end users to correctly apply rules in practice, for instance when sharing documents, sending emails, or collaborating across different platforms and organizations.

Microsoft customers indicate that approximately 50% of an Information Protection project is business focused rather than technical, so end-user training and communication is critical to success.

Resiliency and continuity overview

Article • 06/05/2023

How does Microsoft ensure business continuity if a disaster or other threat to service availability occurs?

Microsoft's Enterprise Resilience and Crisis Management (ERCM) team oversees business continuity management and disaster recovery activities across Microsoft services and cloud offerings. Representatives from Microsoft business units coordinate with the ERCM team to develop business continuity plans and validate compliance with business continuity requirements.

The Business Continuity Management (BCM) lifecycle is at the core of our BCM methodology. This three-phase process is designed to be adaptable so it can be implemented by a wide variety of business models across Microsoft. It begins with an **Assessment** phase to identify critical processes and objectives that should be included in the business continuity program. The Assessment phase also requires a Business Impact Analysis (BIA). The **Planning** phase focuses on developing and implementing resilience and recovery strategies and documenting them in official business continuity plans. Finally, **Capability Validation** tests business continuity plans and their implementations to verify effectiveness and identify potential improvements.

Microsoft online services business continuity strategies use hardware, network, and datacenter redundancy. Data replication between datacenters provides high availability and reliability during a catastrophic incident. It also increases resilience to mundane incidents such as isolated hardware failure or data corruption.

How does Microsoft test business continuity and disaster recovery plans?

Microsoft's Enterprise Resilience and Crisis Management (ERCM) policy stipulates that all Microsoft business continuity and disaster recovery plans must be tested, updated, and reviewed on an annual basis. Microsoft online services test their business continuity plans at least annually per ERCM policies. After Action reports are created and reviewed to validate, test results and inform plan updates in response to any problems discovered during testing.

To validate resilience and recovery strategies against a wide range of potential incidents, the ERCM Program defines multiple categories of test scenarios affecting people, locations, and technology. The level of validation required for each service is based on the service's criticality, with more critical services receiving more rigorous validation. Each Microsoft online service team tests their business continuity plan according to ERCM guidelines to measure the plan's effectiveness and the service team's readiness to execute the plan.

Per ERCM guidelines, annual reviews of business continuity plans and capability validation must take place within 12 months of the last review. Capability validation must include review of supporting documentation, such as the BIA, to ensure it remains accurate. Microsoft makes capability validation results for select Microsoft online services available to our customers through quarterly reports.

How do Microsoft online services ensure system capacity meets demand?

Capacity planning helps service teams allocate the resources necessary to support Microsoft online service availability. Regular capacity planning is required as part of Microsoft's ERCM program. Service teams review capacity data during quarterly reviews, and during emergency situations that warrant more capacity review.

The raw data for capacity planning is maintained by each service team and includes metrics like system processing, memory, and hardware capacity. Scheduled reviews use a model of the system's current capacity and test it against projected needs in emergency situations. If the model indicates gaps in capacity, proposed changes to system capacity are submitted to service team leadership for review. Approved changes are incorporated into a new model before implementation by service team engineers.

How do Microsoft online services maintain service availability during routine system failures?

Microsoft online services achieve service resilience through redundant architecture, data replication, and automated integrity checking. Redundant architecture involves deploying multiple instances of a service on geographically and physically separate hardware, providing increased fault-tolerance for Microsoft online services. Data replication ensures there are always multiple copies of customer data in different fault-zones, allowing critical customer data to be recovered if corrupted, lost, or even

accidentally deleted by the customer. Automated integrity checking increases data availability by automatically restoring data impacted by many kinds of physical or logical corruption.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to resiliency and continuity.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
ISO 27001/27002 Statement of Applicability Certificate	A.17.1: Information security continuity A.17.2: Redundancies	April 24, 2023
ISO 22301 Certificate	All controls	April 24, 2023
SOC 1 SOC 2 SOC 3	BC-1: Business continuity plans BC-3: Business continuity and disaster recovery procedures BC-4: BCDR testing BC-7: Datacenter business continuity plans BC-8: Datacenter business continuity testing BC-9: Datacenter resiliency assessment DS-5: Backup key service components DS-6: Redundancy of critical components DS-7: Automatic replication of customer data DS-8: Backup schedule DS-9: Backup restoration procedures DS-11: Offsite backups DS-14: Automatic restoration of customer services	November 17, 2023

Microsoft 365

Expand table

External audits	Section	Latest report date
FedRAMP (Office 365)	CP-2: Contingency plan CP-3: Contingency training CP-4: Contingency plan testing CP-6: Alternate storage site CP-7: Alternate processing site CP-9: Information system backup CP-10: Information system recovery and reconstitution	July 31, 2023
ISO 27001/27002 Statement of Applicability Certificate	A.17.1: Information security continuity A.17.2: Redundancies	March 2023
ISO 22301 Certificate	All controls	March 2023
SOC 1 SOC 2	CA-49: Backup policies CA-50: Business continuity CA-51: Data replication	January 3, 2023
SOC 3	CUEC-09: EXO email restoration	January 3, 2023

Resources

- Video: Microsoft online services continuity management
- Microsoft Cloud ERCM: Business Continuity and Disaster Recovery Plan Validation Report FY24 Q1
- Enterprise Resilience and Crisis Management (ERCM) Program

Legal disclaimer

- Enterprise business continuity legal disclaimer

Feedback

Was this page helpful? Yes No

Azure infrastructure availability

Article • 01/22/2023

This article provides information about what Microsoft does to secure the Azure infrastructure and provide maximum availability of customers' data. Azure provides robust availability, based on extensive redundancy achieved with virtualization technology.

Temporary outages and natural disaster

The Microsoft Cloud Infrastructure and Operations team designs, builds, operates, and improves the security of the cloud infrastructure. This team ensures that the Azure infrastructure is delivering high availability and reliability, high efficiency, and smart scalability. The team provides a more secure, private, and trusted cloud.

Uninterruptible power supplies and vast banks of batteries ensure that electricity remains continuous if a short-term power disruption occurs. Emergency generators provide backup power for extended outages and planned maintenance. If a natural disaster occurs, the datacenter can use onsite fuel reserves.

High-speed and robust fiber optic networks connect datacenters with other major hubs and internet users. Compute nodes host workloads closer to users to reduce latency, provide geo-redundancy, and increase overall service resiliency. A team of engineers works around the clock to ensure services are persistently available.

Microsoft ensures high availability through advanced monitoring and incident response, service support, and backup failover capability. Geographically distributed Microsoft operations centers operate 24/7/365. The Azure network is one of the largest in the world. The fiber optic and content distribution network connects datacenters and edge nodes to ensure high performance and reliability.

Disaster recovery

Azure keeps your data durable in two locations. You can choose the location of the backup site. In the primary location, Azure constantly maintains three healthy replicas of your data.

Database availability

Azure ensures that a database is internet accessible through an internet gateway with sustained database availability. Monitoring assesses the health and state of the active databases at five-minute time intervals.

Storage availability

Azure delivers storage through a highly scalable and durable storage service, which provides connectivity endpoints. This means that an application can access the storage service directly. The storage service processes incoming storage requests efficiently, with transactional integrity.

Next steps

To learn more about what Microsoft does to help secure the Azure infrastructure, see:

- Azure facilities, premises, and physical security
- Azure information system components and boundaries
- Azure network architecture
- Azure production network
- Azure SQL Database security features
- Azure production operations and management
- Azure infrastructure monitoring
- Azure infrastructure integrity
- Azure customer data protection

Azure reliability documentation

Article • 09/24/2023

Reliability consists of two principles: resiliency and availability. The goal of resiliency is to avoid failures and, if they still occur, to return your application to a fully functioning state. The goal of availability is to provide consistent access to your application or workload. It's important to plan for proactive reliability based on your application requirements.

Azure includes built-in reliability services that you can use and manage based on your business needs. Whether it's a single hardware node failure, a rack level failure, a datacenter outage, or a large-scale regional outage, Azure provides solutions that improve reliability. For example, availability sets ensure that the virtual machines deployed on Azure are distributed across multiple isolated hardware nodes in a cluster. Availability zones protect customers' applications and data from datacenter failures across multiple physical locations within a region. **Regions** and **availability zones** are central to your application design and resiliency strategy and are discussed in greater detail later in this article.

The Azure reliability documentation offers reliability guidance for Azure services. This guidance includes information on availability zone support, disaster recovery guidance, and availability of services.

For detailed service-specific reliability guidance, including availability zones, disaster recovery, or high availability, see Azure service-specific reliability guidance overview.

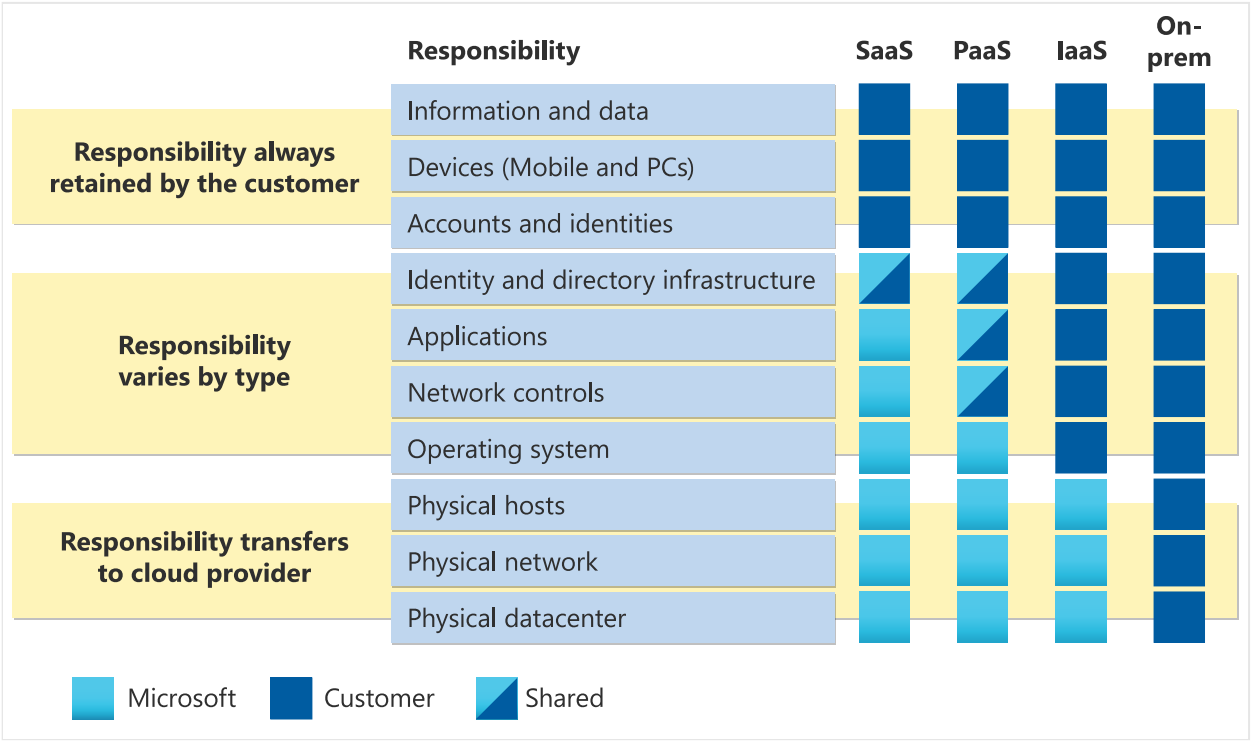
For information on reliability and reliability principles and architecture in Microsoft Azure services, see Microsoft Azure Well-Architected Framework: Reliability.

Reliability requirements

The required level of reliability for any Azure solution depends on several considerations. Availability and latency SLA and other business requirements drive the architectural choices and resiliency level and should be considered first. Availability requirements range from how much downtime is acceptable – and how much it costs your business – to the amount of money and time that you can realistically invest in making an application highly available.

Building reliability systems on Azure is a **shared responsibility**. Microsoft is responsible for the reliability of the cloud platform, including its global network and data centers. Azure customers and partners are responsible for the resilience of their cloud

applications, using architectural best practices based on the requirements of each workload. While Azure continually strives for highest possible resiliency in SLA for the cloud platform, you must define your own target SLAs for each workload in your solution. An SLA makes it possible to evaluate whether the architecture meets the business requirements. As you strive for higher percentages of SLA guaranteed uptime, the cost and complexity to achieve that level of availability grows. An uptime of 99.99 percent translates to about five minutes of total downtime per month. Is it worth the more complexity and cost to reach that percentage? The answer depends on the individual business requirements. While deciding final SLA commitments, understand Microsoft's supported SLAs. Each Azure service has its own SLA.



In the traditional on-premises model, the entire responsibility of managing, from the hardware for compute, storage and networking to the application, falls on you. You must plan for various types of failures and how to deal with them by creating a disaster recovery strategy. With IaaS, the cloud service provider is responsible for the core infrastructure resiliency, including storage, networking, and compute. As you move from IaaS to PaaS and then to SaaS, you'll find that you're responsible for less and the cloud service provider is responsible for more.

For more information on Reliability principles, see Well-architected Framework Reliability documentation.

Building reliability

You should define your application's reliability requirements at the beginning of planning. If you know which applications don't need 100% high availability during

certain periods of time, you can optimize costs during those non-critical periods. Identify the type of failures an application can experience, and the potential effect of each failure. A recovery plan should cover all critical services by finalizing recovery strategy at the individual component and the overall application level. Design your recovery strategy to protect against zonal, regional, and application-level failure. Perform testing of the end-to-end application environment to measure application reliability and recovery against unexpected failure.

The following checklist covers the scope of reliability planning.

Expand table

Reliability planning
Define availability and recovery targets to meet business requirements.
Design the reliability features of your applications based on the availability requirements.
Align applications and data platforms to meet your reliability requirements.
Configure connection paths to promote availability.
Use availability zones and disaster recovery planning where applicable to improve reliability and optimize costs.
Ensure your application architecture is resilient to failures.
Know what happens if SLA requirements are not met.
Identify possible failure points in the system; application design should tolerate dependency failures by deploying circuit breaking.
Build applications that operate in the absence of their dependencies.

RTO and RPO

Two important metrics to consider are the recovery time objective and recovery point objective, as they pertain to disaster recovery. For more information on functional and non-functional design requirements, see Well-architected Framework functional and nonfunctional requirements.

- **Recovery time objective (RTO)** is the maximum acceptable time that an application can be unavailable after an incident.
- **Recovery point objective (RPO)** is the maximum duration of data loss that is acceptable during a disaster.

RTO and RPO are non-functional requirements of a system and should be dictated by business requirements. To derive these values, it's a good idea to conduct a risk assessment, and clearly understanding the cost of downtime or data loss.

Regions and availability zones

Regions and availability zones are a big part of the reliability equation. Regions feature multiple, physically separate availability zones. These availability zones are connected by a high-performance network featuring less than 2ms latency between physical zones. Low latency helps your data stay synchronized and accessible when things go wrong. You can use this infrastructure strategically as you architect applications and data infrastructure that automatically replicate and deliver uninterrupted services between zones and across regions.

Microsoft Azure services support availability zones and are enabled to drive your cloud operations at optimum high availability while supporting your cross-region recovery and business continuity strategy needs.

For disaster recovery planning, regions that are paired with other regions offer cross-region replication and provide protection by asynchronously replicating data across other Azure regions. Regions without a pair follow data residency guidelines and offer high availability with availability zones and locally redundant or zone-redundant storage. Customers will need to plan for their cross-region disaster recovery based on their RTO/RPO needs.

Choose the best region for your needs based on technical and regulatory considerations —service capabilities, data residency, compliance requirements, latency—and begin advancing your reliability strategy. For more information, see Azure regions and availability zones.

Azure service dependencies

Microsoft Azure services are available globally to drive your cloud operations at an optimal level.

Azure services deployed to Azure regions are listed on the Azure global infrastructure products page. To better understand regions and Availability Zones in Azure, see Regions and Availability Zones in Azure.

Azure services are built for reliability including high availability and disaster recovery. There are no services that are dependent on a single logical data center (to avoid single points of failure). Non-regional services listed on Azure global infrastructure products

are services for which there is no dependency on a specific Azure region. Non-regional services are deployed to two or more regions and if there is a regional failure, the instance of the service in another region continues servicing customers. Certain non-regional services enable customers to specify the region where the underlying virtual machine (VM) on which service runs will be deployed. For example, Azure Virtual Desktop enables customers to specify the region location where the VM resides. All Azure services that store customer data allow the customer to specify the specific regions in which their data will be stored. The exception is Microsoft Entra ID, which has geo placement (such as Europe or North America). For more information about data storage residency, see the Data residency map.

If you need to understand dependencies between Azure services to help better architect your applications and services, you can request the **Azure service dependency documentation** by contacting your Microsoft sales or customer representative. This document lists the dependencies for Azure services, including dependencies on any common major internal services such as control plane services. To obtain this documentation, you must be a Microsoft customer and have the appropriate non-disclosure agreement (NDA) with Microsoft.

Next steps

- For service specific guides on availability zone support and disaster recovery, see Reliability guidance.
- For service migration guides to availability zone support, see Availability zone migration guidance.
- Business continuity management in Azure
- Availability of service by category
- Build solutions for high availability using availability zones
- What are Azure regions and availability zones?
- Cross-region replication in Azure | Microsoft Learn
- Training: Describe high availability and disaster recovery strategies

Azure network architecture

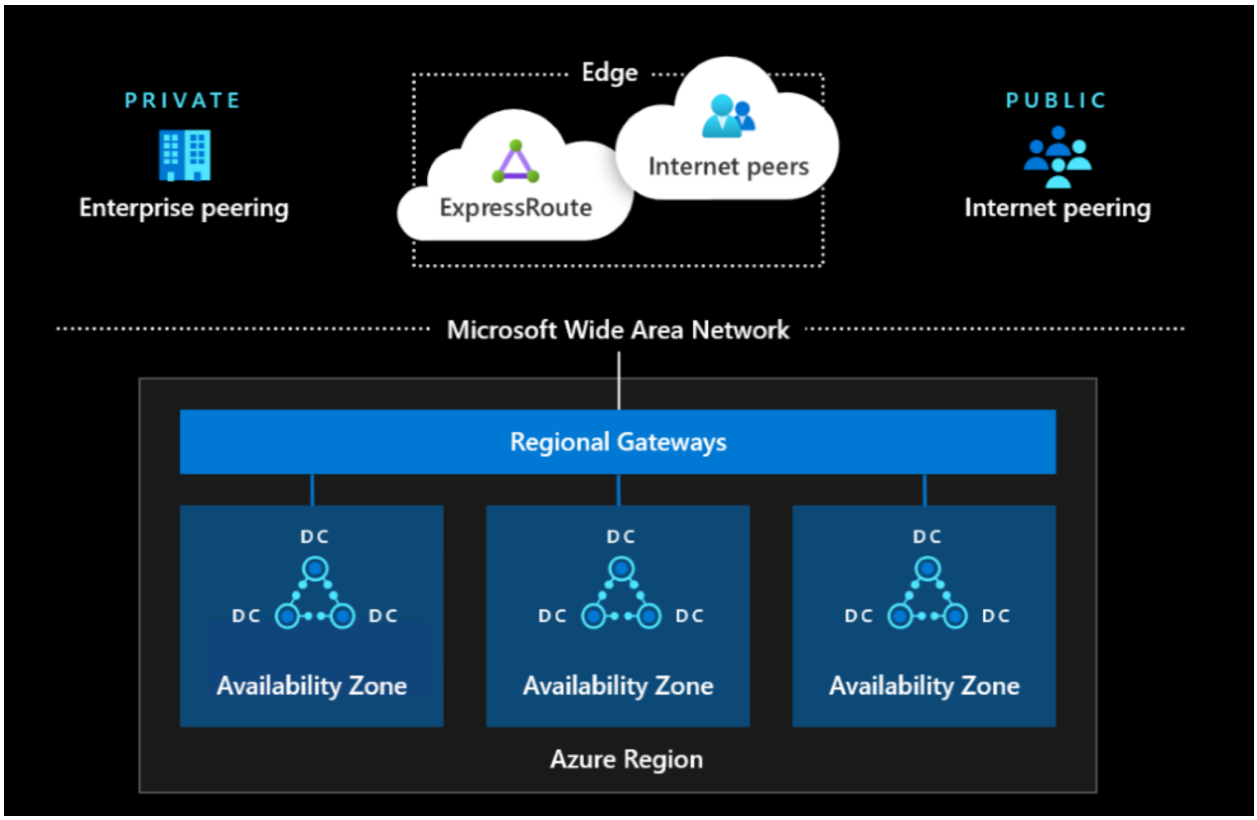
Article • 02/28/2023

The Azure network architecture provides connectivity from the Internet to the Azure datacenters. Any workload deployed (IaaS, PaaS, and SaaS) on Azure is leveraging the Azure datacenter network.

Network topology

The network architecture of an Azure datacenter consists of the following components:

- Edge network
- Wide area network
- Regional gateways network
- Datacenter network



Network components

A brief description of the network components.

- Edge network
 - Demarcation point between Microsoft networking and other networks (for example, Internet, Enterprise network)

 - Provides Internet and ExpressRoute peering into Azure

- Wide area network
 - Microsoft intelligent backbone network covering the globe
 - Provides connectivity between Azure regions

- Regional gateway
 - Point of aggregation for all of the datacenters in an Azure region
 - Provides massive connectivity between datacenters within an Azure region (for example, multi hundred terabits per datacenter)

- Datacenter network
 - Provides connectivity between servers within the datacenter with low oversubscribed bandwidth

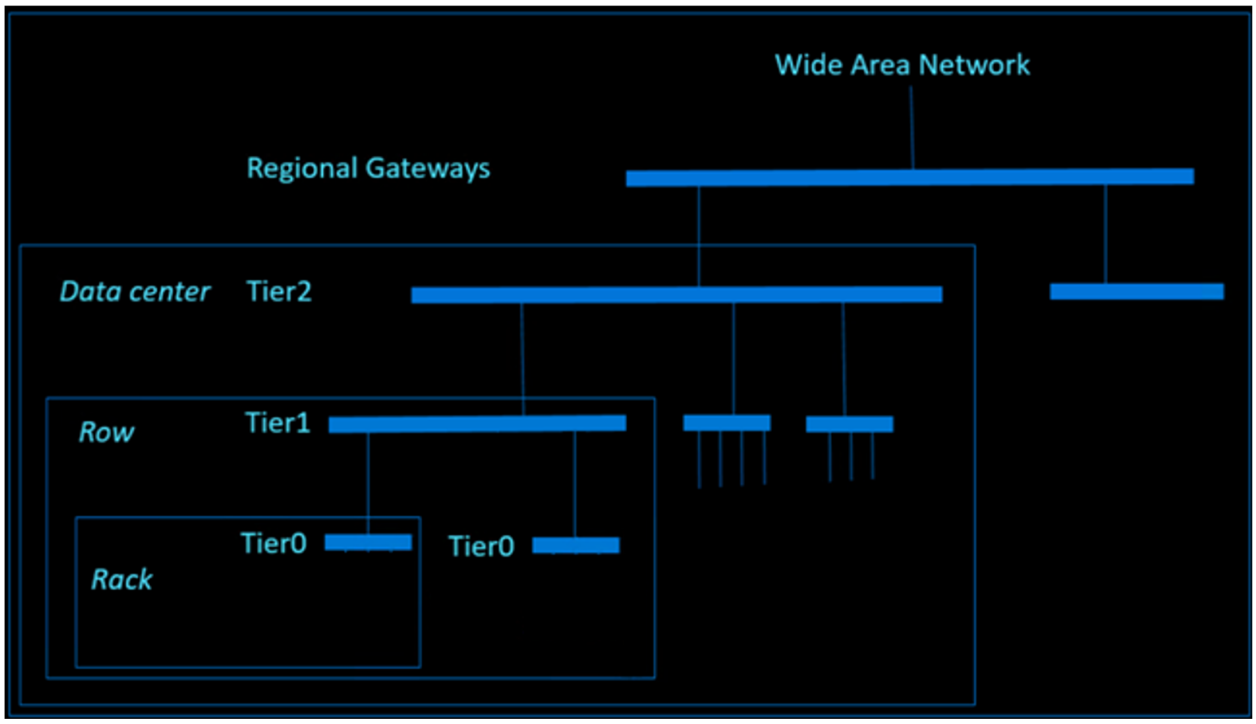
The above network components are designed to provide maximum availability to support always-on, always-available cloud business. The redundancy is designed and built into the network from the physical aspect all the way up to control protocol.

Datacenter network resiliency

Let's illustrate the resiliency design principle using datacenter network.

The datacenter network is a modified version of a Clos network, providing high bi-sectional bandwidth for cloud scale traffic. The network is constructed using a large number of commodity devices to reduce the impact caused by individual hardware failure. These devices are strategically located in different physical locations with separate power and cooling domain to reduce impact of an environment event. On the control plane, all network devices are running as OSI model Layer 3 routing mode, which eliminates the historical issue of traffic loop. All paths between different tiers are active to provide high redundancy and bandwidth using Equal-Cost Multi-Path (ECMP) Routing.

The following diagram demonstrates that the datacenter network is constructed by different tiers of network devices. The bars in the diagram represent groups of network devices which provide redundancy and high bandwidth connectivity.



Next steps

To learn more about what Microsoft does to help secure the Azure infrastructure, see:

- Azure facilities, premises, and physical security
- Azure infrastructure availability
- Azure information system components and boundaries
- Azure production network
- Azure SQL Database security features
- Azure production operations and management
- Azure infrastructure monitoring
- Azure infrastructure integrity
- Azure customer data protection

Exchange Online data resiliency in Microsoft 365

Article • 03/02/2023

> **ⓘ Important**
>
> As we continue to invest in different ways to preserve mailbox content, we're announcing the retirement of In-Place Holds in the Exchange admin center (EAC) in Exchange Online. Starting July 1, 2020, you won't be able to create new In-Place Holds. But you'll still be able to manage In-Place Holds in the EAC or by using the **Set-MailboxSearch** cmdlet in Exchange Online PowerShell. However, starting October 1, 2020, you won't be able to manage In-Place Holds. You'll only be able to remove them in the EAC or by using the **Remove-MailboxSearch** cmdlet. Using In-Place Holds in Exchange Server and Exchange hybrid deployments will still be supported. For more information about the retirement of In-Place Holds in Exchange Online, see **Retirement of legacy eDiscovery tools**.

An In-Place Hold preserves all mailbox content, including deleted items and original versions of modified items. All such mailbox items are returned in an In-Place eDiscovery search. When you place an In-Place Hold on a user's mailbox, the contents in the corresponding archive mailbox (if it's enabled) are also placed on hold and returned in an eDiscovery search.

There are two types of corruption that can affect an Exchange database: physical corruption, which is typically caused by hardware (in particular, storage hardware) problems, and logical corruption, which occurs due to other factors. Generally, there are two types of logical corruption that can occur within an Exchange database:

- **Database logical corruption** - The database page checksum matches, but the data on the page is wrong logically. This can occur when the database engine (the Extensible Storage Engine (ESE)) attempts to write a database page and even though the operating system returns a success message, the data is either never written to the disk or it's written to the wrong place. This is referred to as a *lost flush*. ESE includes numerous features and safeguards that are designed to prevent physical corruption of a database and other data loss scenarios. To prevent lost flushes from losing data, ESE includes a lost flush detection mechanism in the database along with a feature (single page restore) to correct it.
- **Store logical corruption** - Data is added, deleted, or manipulated in a way that the user doesn't expect. These cases are caused by third-party applications. It's usually

corruption in the sense that the user views it as corruption. The Exchange store considers the transaction that produced the logical corruption to be a series of valid MAPI operations. The In-Place Hold features in Exchange Online provides protection from store logical corruption (because it prevents content from being permanently deleted by a user or an application).

Exchange Online performs several consistency checks on replicated log files during both log inspection and log replay. These consistency checks prevent physical corruption from being replicated by the system. For example, during log inspection, there is a physical integrity check that verifies the log file and validates that the checksum recorded in the log file matches the checksum generated in memory. In addition, the log file header is examined to make sure the log file signature recorded in the log header matches that of the log file. During log replay, the log file undergoes further scrutiny. For example, the database header also contains the log signature that is compared with the log file's signature to ensure they match.

Protection against corruption of mailbox data in Exchange Online is achieved by using Exchange Native Data Protection, a resiliency strategy that leverages application-level replication across multiple servers and multiple datacenters along with other features that help protect data from being lost due to corruption or other reasons. These features include native features that are managed by Microsoft or the Exchange Online application itself, such as:

- Data Availability Groups
- Single Bit Correction
- Online Database Scanning
- Lost Flush Detection
- Single Page Restore
- Mailbox Replication Service
- Log File Checks
- Deployment on Resilient File System

For more information on the native features listed previously, select the hyperlinks, and see the following for additional information and for details on items without hyperlinks. In addition to these native features, Exchange Online also includes data resiliency features that customers can manage, such as:

- Single Item Recovery (enabled by default)
- In-Place Hold and Litigation Hold
- Deleted Item Retention and Soft-Deleted Mailboxes (both enabled by default)

Database Availability Groups

Every mailbox database in Microsoft 365 is hosted in a database availability group (DAG) and replicated to geographically separate datacenters within the same region. The most common configuration is four database copies in four datacenters; however, some regions have fewer datacenters (databases are replicated to three datacenters in India, and two datacenters in Australia and Japan). But in all cases, every mailbox database has four copies that are distributed across multiple datacenters, thereby ensuring that mailbox data is protected from software, hardware, and even datacenter failures.

Out of these four copies, three of them are configured as highly available. The fourth copy is configured as a lagged database copy. The lagged database copy is not intended for individual mailbox recovery or mailbox item recovery. Its purpose is to provide a recovery mechanism for the rare event of system-wide, catastrophic logical corruption.

Lagged database copies in Exchange Online are configured with a seven-day log file replay lag time. In addition, the Exchange Replay Lag Manager is enabled to provide dynamic log file play down for lagged copies to allow lagged database copies to self-repair and manage log file growth. Although lagged database copies are used in Exchange Online, it is important to understand that they are not a guaranteed point-in-time backup. Lagged database copies in Exchange Online have an availability threshold, typically around 90%, due to periods where the disk containing a lagged copy is lost due to disk failure, the lagged copy becoming a highly available copy (due to automatic play down), as well as the periods where the lagged database copy is rebuilding the log replay queue.

Transport Resilience

Exchange Online includes two primary transport resilience features: Shadow Redundancy and Safety Net. Shadow Redundancy keeps a redundant copy of a message while it is in transit. Safety Net keeps a redundant copy of a message after the message is successfully delivered.

With Shadow Redundancy, each Exchange Online transport server makes a copy of each message it receives before it acknowledges successfully receiving the message to the sending server. This makes all messages in the transport pipeline redundant while in transit. If Exchange Online determines the original message was lost in transit, a redundant copy of the message is redelivered.

Safety Net is a transport queue that is associated with the Transport service on a Mailbox server. This queue stores copies of messages that were successfully processed by the server. When a mailbox database or server failure requires activating an out-of-date copy of the mailbox database, messages in the Safety Net queue are automatically

resubmitted to the new active copy of the mailbox database. Safety Net is also redundant, thereby eliminating transport as a single point of failure. It uses the concept of a Primary Safety Net and a Shadow Safety Net wherein if the Primary Safety Net is unavailable for more than 12 hours, resubmit requests become shadow resubmit requests, and messages are redelivered from the Shadow Safety Net.

Message resubmissions from Safety Net are automatically initiated by the Active Manager component of the Microsoft Exchange Replication service that manages DAGs and mailbox database copies. No manual actions are required to resubmit messages from Safety Net.

Single Bit Correction

ESE includes a mechanism to detect and resolve single-bit CRC errors (also known as single-bit flips) that are the result of hardware errors (and as such they represent physical corruption). When these errors occur, ESE automatically corrects them and logs an event in the event log.

Online Database Scanning

Online database scanning (also known as *database check summing*) is the process where an ESE uses a database consistency checker to read each page and check for page corruption. The primary purpose is to detect physical corruption and lost flushes that may not be getting detected by transactional operations. Database scanning also performs post-store crash operations. Space can be leaked due to crashes, and online database scanning finds and recovers lost space. The system is designed with the expectation that every database is fully scanned once every seven days.

Lost Flush Detection

A lost flush occurs when a database write operation that the disk subsystem/operating system returned as completed did not actually get written to disk, or was written in the wrong location. Lost flush incidents can result in database logical corruption, so to prevent lost flushes from resulting in lost data, ESE includes a lost flush detection mechanism. As database pages are written to passive copies, a check is performed for lost flushes on the active copy. If a lost flush is detected, ESE can repair the process using a page patching process.

Single Page Restore

Single page restore, also known as *page patching*, is an automatic process where corrupt database pages are replaced by healthy copies from a healthy replica. The repair process for a corrupt page depends on whether the database copy is active or passive. When an active database copy encounters a corrupted page, it can copy a page from one of its replicas, provided the page it copies is up to date. This process is accomplished by putting a request for the page into the log stream, which is the basis of mailbox database replication. As soon as a replica encounters the page request, it responds by sending a copy of the page to the requesting database copy. Single page restore also provides an asynchronous communication mechanism for the active to request a page from replicas, even if the replicas are currently offline.

In case of corruption in a passive database copy, including a lagged database copy, because these copies are always behind their active copy, it is always safe to copy any page from the active copy to a passive copy. A passive database copy is by nature highly available, so during the page patching process, log replaying is suspended, but log copying continues. The passive database copy retrieves a copy of the corrupted page from the active copy, waits until the log file that meets the maximum required log generation requirement is copied and inspected, and then patches the corrupt page. Once the page has been patched, log replay resumes. The process is the same for the lagged database copy, except that the lagged database first replays all log files that are necessary to achieve a patchable state.

Mailbox Replication Service

Moving mailboxes is a key part of managing a large-scale email service. There are always updated technologies and hardware and version upgrades to deal with, so having a robust, throttled system that enables our engineers to accomplish this work while keeping the mailbox moves transparent to users (by making sure they stay online throughout the process) is key and making sure that the process scales up gracefully as mailboxes get larger and larger.

The Exchange Mailbox Replication Service (MRS) is responsible for moving mailboxes between databases. During the move, MRS performs a consistency check on all items within the mailbox. If a consistency issue is found, MRS will either correct the problem, or skip the corrupted items, thereby removing the corruption from the mailbox.

Because MRS is a component of Exchange Online, we can make changes in its code to address new forms of corruption that are detected in the future. For example, if we detect a consistency issue that MRS is not able to fix, we can analyze the corruption, change the MRS code and correct the inconsistency (if we understand how to).

Log File Checks

All transaction log files generated by an Exchange database undergo several forms of consistency checks. When a log file is created, the first thing done is a bit pattern is written and then a series of log writes is performed. This structure enables Exchange Online to execute a series of checks (lost flush, CRC, and other checks) to validate each log file as it is written, and again as it is replicated.

Deployment on Resilient File System

To help prevent corruption from occurring at the file system level, Exchange Online is being deployed on Resilient File System (ReFS) partitions to provide improved recovery capabilities. ReFS is a file system in Windows Server 2012 and later that is designed to be more resilient against data corruption thereby maximizing data availability and integrity. Specifically, ReFS brings improvements in the way that metadata is updated which offers better protection for data and reduces data corruption cases. It also uses checksums to verify the integrity of file data and metadata ensuring that data corruption is easily found and repaired.

Exchange Online takes advantage of several ReFS benefits:

- More resiliency in data integrity means fewer data corruption incidents. Reducing the number of corruption incidents means fewer unnecessary database reseeds.
- Checksum running on metadata enabling detections of corruption cases sooner and more deterministically, allowing us to fix customer data corruption before grey failures occur on data volumes.
- Designed to work well with large data sets—petabytes and larger—without performance impact
- Support for other features used by Exchange Online, such as BitLocker encryption.

Exchange Online also benefits from other ReFS features:

- **Integrity (Integrity Streams)** - ReFS stores data in a way that protects it from many of the common errors that can normally cause data loss. Microsoft 365 Search uses Integrity Streams to help with early disk corruption detection and checksums of file content. The feature also reduces corruption incidents caused by 'Torn Writes' (when a write operation does not complete due to power outages, etc.).
- **Availability (Salvage)** - ReFS prioritizes the availability of data. Historically, file systems were often susceptible to data corruption that would require the system to be taken offline for repair. Although rare, if corruption does occur, ReFS implements salvage, a feature that removes the corrupt data from the namespace on a live volume and ensures that good data is not adversely affected by non-

repairable corrupt data. Applying the Salvage feature and isolating data corruption to Exchange Online database volumes means that we can keep non-affected databases on a corrupted volume healthy between the time of corruption and repair action. This structure increases the availability of databases that would normally be affected by such disk corruption issues.

SharePoint and OneDrive data resiliency in Microsoft 365

Article • 03/02/2023

Within Microsoft 365, OneDrive is built on top of the SharePoint file platform. In this article, only SharePoint will be used to refer to both products. The contents of this article are relevant to Microsoft 365 and don't apply to consumer services.

There are two primary assets that make up the core content storage of SharePoint:

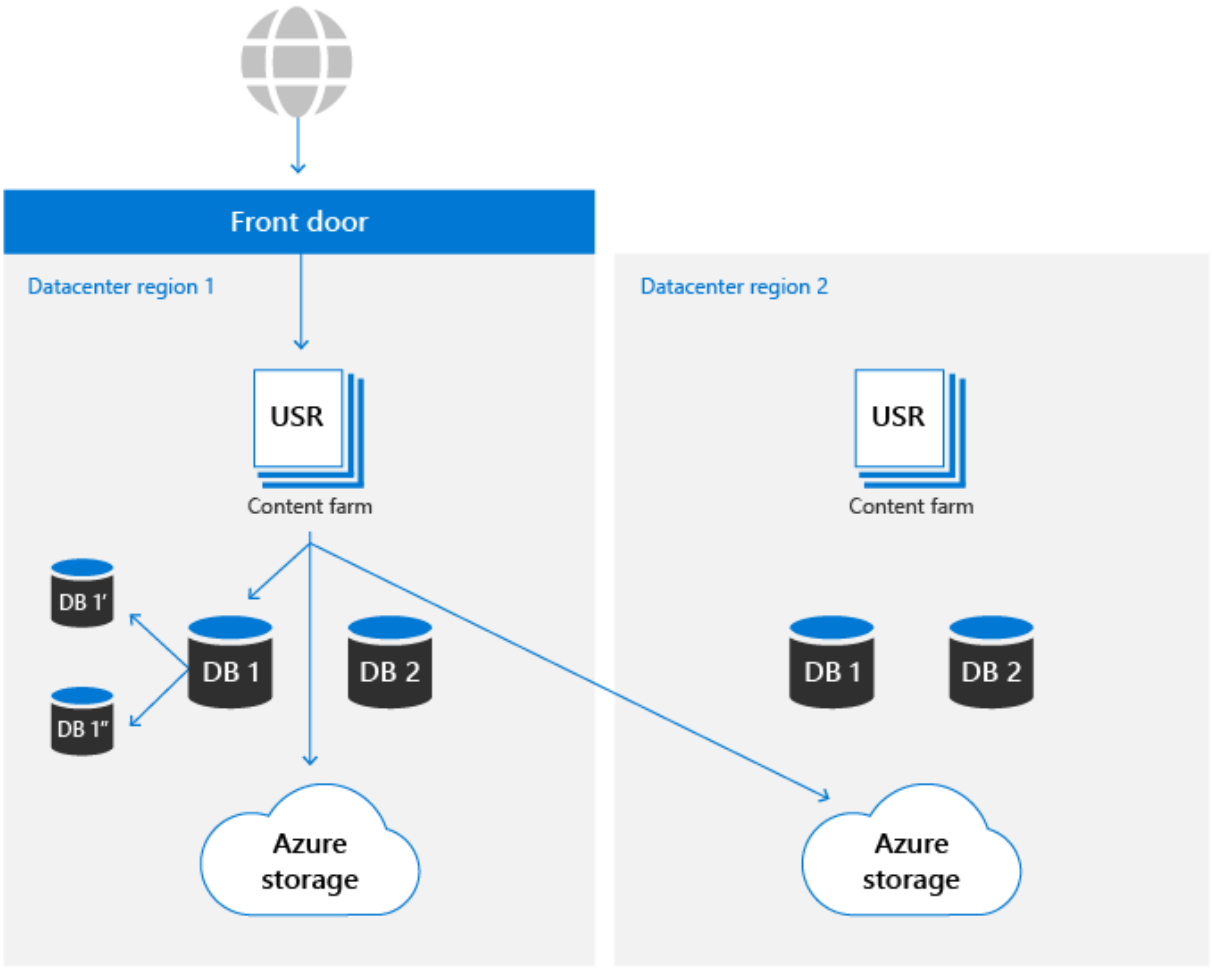
- **Metadata**: Metadata about each file is stored in Azure SQL Database. Azure SQL offers a complete business continuity story that SharePoint uses and details are covered later in this article.
- **Blob storage**: User content that is uploaded into SharePoint is stored in Azure Storage. SharePoint has built a custom resiliency plan on top of Azure Storage to ensure near real-time duplication of user content and a truly active/active system.

The complete set of controls to ensure data resiliency is explained in further sections.

Blob storage resilience

SharePoint has a custom-built solution for storage of customer data in Azure Storage. Every file is simultaneously written into both a primary and a secondary datacenter region. If writes to either Azure region fail, the file save will fail. After the contents are written into Azure Storage, checksums are stored separately with metadata, and are used to ensure that the committed write is identical to the original file sent to SharePoint during all future reads. This same technique is used in all workflows to prevent propagation of any corruption that should occur. Within each region, Azure Locally Redundant Storage (LRS) provides a high level of reliability. See the Azure Storage redundancy article for details.

SharePoint uses Append-Only storage. This means that Microsoft can only add new blobs and can never change old ones until they're permanently deleted. This process ensures that files can't be changed or corrupted after an initial save, protecting against attackers that try to corrupt old versions. Since version integrity protection is built into SharePoint's architecture, previous versions of the file contents can be retrieved, depending on individual administrator settings.

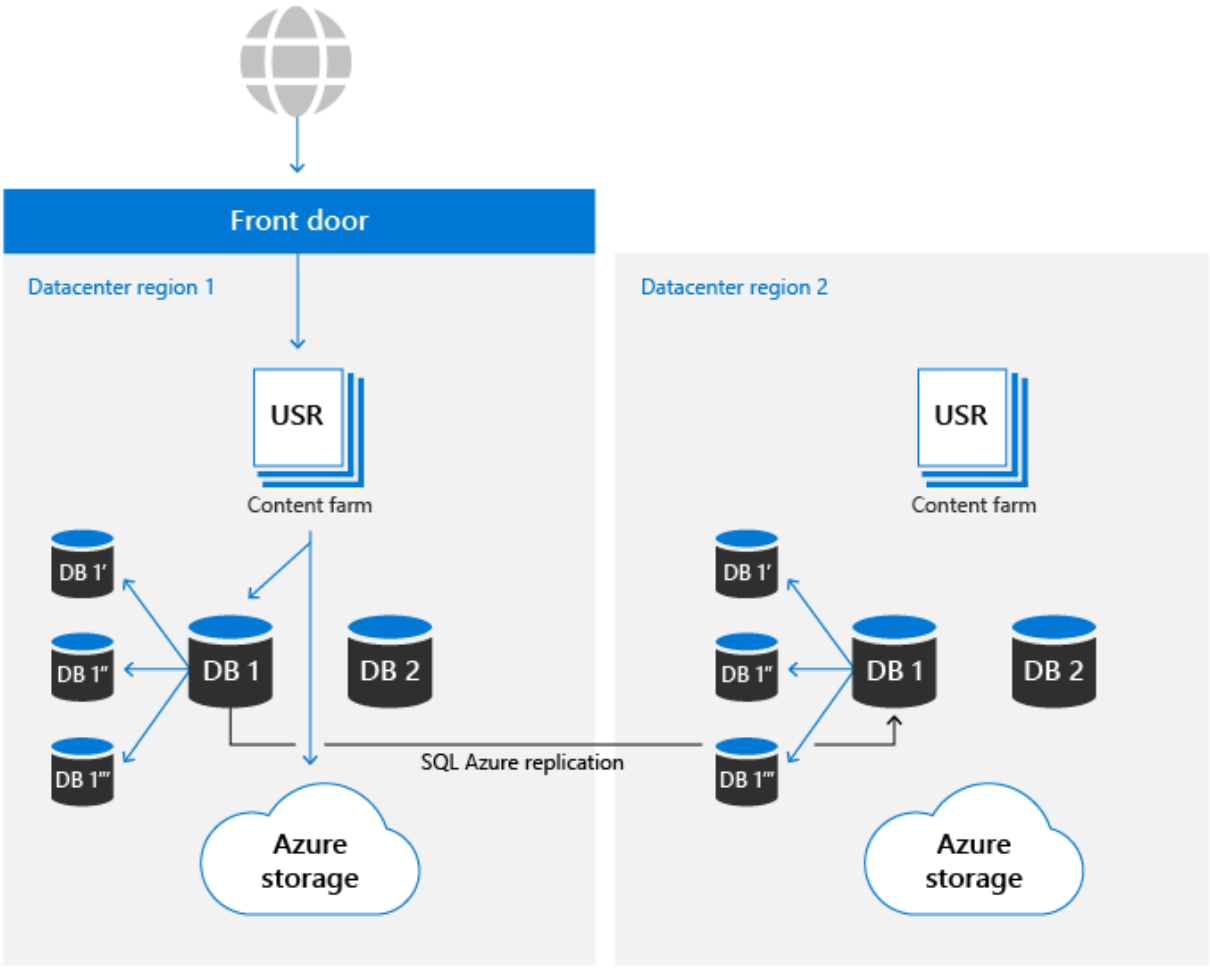


SharePoint environments in either datacenter can access storage containers in both Azure regions. For performance reasons the storage container in the same local datacenter is always preferred, however, read requests that don't see results within a desired threshold will have the same content requested from the remote datacenter to ensure data is always available.

Metadata resilience

SharePoint metadata is also critical to accessing user content as it stores the location of and access keys to the content stored in Azure Storage. These databases are stored in Azure SQL, which has an extensive business continuity plan.

SharePoint uses the replication model provided by Azure SQL and has built a proprietary automation technology to determine a failover is required and to initiate the operation if necessary. As such, it falls into the 'Manual database failover' category from an Azure SQL perspective. The latest metrics for Azure SQL database recoverability are available here.



SharePoint uses Azure SQL's backup system to enable Point in Time Restores (PITR) for up to 14 days. PITR is covered more in a later section.

Automated failover

SharePoint uses a custom-built automated failover to minimize the impact on customer experience when a location-specific event occurs. Monitoring-driven automation detecting a single or multi-component failure beyond certain thresholds will result in automated redirection of all users' activity out of the problematic environment and over to a warm secondary. A failover results in metadata and compute storage being served entirely out of the new datacenter. As blob storage always runs entirely active/active, no change is required for a failover. The compute tier will prefer the nearest blob container but will use both local and remote blob storage locations at any time to ensure availability.

SharePoint uses the Azure Front Door service to provide routing internal to the Microsoft network. This configuration allows failover redirection independent of DNS and reduces the effect of local-machine caching. Most failover operations are

transparent to end users. If there's a failover, customers won't need to make any changes in order to maintain access to the service.

Versioning and Files Restore

For newly created document libraries, SharePoint defaults to 500 versions on every file and can be configured to retain more versions if desired. The UI doesn't allow a value fewer than 100 versions to be set, but it's possible to set the system to store fewer versions using public APIs. For reliability, any value less than 100 isn't recommended and can result in user activity causing inadvertent data loss.

For more information about versioning, see Versioning in SharePoint.

Files Restore is the ability to go 'back in time' on any Document Library in SharePoint to any second of time in the last 30 days. This process can be used to recover from ransomware, mass deletions, corruption, or any other event. This feature uses file versions so reducing default versions can reduce the effectiveness of this restore.

The Files Restore feature is documented for both OneDrive and SharePoint.

Deletion, backup, and Point in Time Restore

User content deleted from SharePoint goes through the following deletion flow.

Deleted items are retained in recycle bins for a certain period of time. For SharePoint, the retention time is 93 days. It begins when you delete the item from its original location. When you delete the item from the site recycle bin, it goes into the site collection recycle bin. It stays there for the remainder of the 93 days, and then is permanently deleted. More information about how to use the recycle bin is available at these links:

- Restore items in the Recycle Bin
- Restore deleted items from the Site Collection Recycle Bin.

This process is the default deletion flow and doesn't take into account retention policies or labels. For more information, see Learn about retention for SharePoint and OneDrive.

After the 93-day recycle pipeline is complete, deletion takes place independently for Metadata and for Blob Storage. Metadata will be removed immediately from the database, which makes the content unreadable unless the metadata is restored from backup. SharePoint maintains 14 days-worth of backups of metadata. These backups are taken locally in near real time and then pushed to storage in redundant Azure Storage

containers on, according to documentation at the time of this publication, a 5-10-minute schedule.

When deleting Blob Storage content, SharePoint utilizes the soft delete feature for Azure Blob Storage to protect against accidental or malicious deletion. With this feature, there are a total of 14 days in which to restore content before it's permanently deleted. Also, because blobs are immutable, Microsoft can always restore the state of a file for a 14 day period.

> ⓘ **Note**
>
> While Microsoft applications will send content to the recycle bin for the standard process, SharePoint does provide APIs that allow for skipping the recycle bin and forcing an immediate delete. Review your applications to ensure this is only done when necessary for compliance reasons.

Integrity Checks

SharePoint uses various methods to ensure the integrity of blobs and metadata at all stages of the data lifecycle:

- **File hash stored in metadata**: Hash of the entire file is stored with file metadata to ensure document level data integrity is maintained during all operations
- **Blob hash stored in metadata**: Each blob-item stores a hash of the encrypted content to protect against corruption in underlying Azure storage.
- **Data integrity job**: Every 14 days, each site is scanned for integrity by listing items in the database and matching those up with listed blobs in Azure storage. The job reports any blob-references missing storage-blobs and can retrieve those blobs through the Azure storage soft-delete feature if needed.

Service resilience in Microsoft Teams

Article • 03/02/2023

Microsoft Teams is engineered from the ground up for the cloud; it has taken a cloud-first, mobile-first, cross-platform approach since inception. Teams is built with a highly resilient microservice architecture that splits tasks across multiple geo-redundant systems to ensure high availability and performance around the world. Microsoft Teams also leverages Azure's built-in resiliency, making it extremely resistant to failures.

Service configuration

In geographies where Microsoft Teams is available, it's operated out of at least two geographically distant Azure regions, and more where possible. The service is run in an Active/Active/n configuration in which users are actively served from all available regions in the geography. Capacity management and planning is such that even in a geography running only the minimum two regions, each is capable of handling the full service load from the entire geography in the event that one Azure region is unavailable.

The microservice architecture means that the Teams service is more reliable, as microservices run in Active/Active state but don't need to be located all within the same Azure regions at the same time. This allows for us to leverage capacity where it's available within a given geography, so that if a site becomes partially unavailable, we can migrate some microservices, without being forced to migrate all of them. Microservices also leverage Azure's native resiliency features.

Failover

Failover operations are used as part of our normal day-to-day service management processes for load balancing and maintenance tasks, as well as for recovery during Service Incidents, so these actions are exercised regularly. Failover decisions are made by the incident management team if deemed appropriate, and automation is leveraged to carry out operations at scale with speed. We're constantly working to improve failover procedures based on learnings from incident postmortems, and in the future are looking at ways to implement automated failover decisions based on machine learning.

Limiting impact of service incidents

Fault isolation

Microsoft Teams service infrastructure is highly segmented, which severely limits the scope of most types of impact. Teams is composed of over 140 microservices by design, in part to limit the impact of single component degradation. As such, most service incidents are limited to only a single component, or a very small subset of users. This is true even of DNS incidents, as many Teams micro-services rely not only on Azure Traffic Manager (ATM), but also third-party (Akamai) for DNS and load balancing.

Graceful degradation

Microsoft Teams has introduced several features to limit the real-world impact of service degradation, including that caused by third-party providers such as ISPs or even a customer's own internal network. Examples include:

- Keeping the media path short – Teams utilizes intelligent routing to reduce the number of network hops and network traversal distance, to reduce the number of potential failure points.
- Handling degradation gracefully – Teams has in-app user notifications when things aren't working great, and suggestions such as turning off video to improve performance if the connection is substandard. Teams will also prompt a user to mute audio if their connection is causing feedback or noise on the line.
- Temporary network loss – Teams will automatically reconnect to a meeting or call once a network connection is restored.
- Fallback to PSTN - Teams will automatically give you the option to rejoin via Public Switched Telephone Network (PSTN) if you're dropped from a call due to a bad network connection. This feature has the added benefit of requiring no conference numbers or passcodes, so it's even easier for users on the go to stay connected.
- Device fail & switch – Teams will detect if an audio device is disconnected and attempt to switch to an alternate device if available.

Code deployment

Teams uses stepped rings of deployment when updating service-side components, starting with small groups of internal test users, followed by the wider Microsoft corporate network. At each step, engineers collect feedback on features and performance, and make every effort to detect any issues prior to deploying in production. When deploying into production, we start with a small subset of systems and then begin deploying in ever-widening rings until all users are connecting to systems running the latest code. Again, any reports of problems related to the update will trigger incident response, and may result in halting an update, rolling back to the last version, and removing degraded systems from active rotation until the problem is

solved. Teams also has the capability to disable specific code paths if necessary, which allows us to very quickly identify and disable problematic code and return the system to a healthy state without having to wait for a new update to roll out.

Dealing with data corruption in Microsoft 365

Article • 04/19/2023

One of the challenging aspects of running a large-scale cloud service is how to handle data corruption, given the large volume of data and independent systems. Data corruption can be caused by:

- Application or infrastructure bugs, corrupting some or all of the application state
- Hardware issues that result in lost data or an inability to read data
- Human operational errors
- Malicious hackers and disgruntled employees
- Incidents in external services that result in some loss of data

Because greater resiliency in data integrity means fewer data corruption incidents, Microsoft has built into Microsoft 365 protection mechanisms to prevent corruption from happening, as well as systems and processes that enable us to recover data if it does. Checks and processes exist within the various stages of the engineering release process to increase resiliency against data corruption, including:

- System Design
- Code organization and structure
- Code review
- Unit tests, integration tests, and system tests
- Trip wires tests/gates

Within Microsoft 365 production environments, peer replication between datacenters ensures that there are always multiple live copies of any data. Standard images and scripts are used to recover lost servers, and replicated data is used to restore customer data. In Exchange Online, every mailbox is hosted in Database Availability Groups (DAGs) and replicated to geographically separate datacenters within the same region. Each mailbox database has four copies distributed across datacenters within the DAG: one active copy, two up-to-date copy, and one 7-day lagged copy used in the rare event of catastrophic logical corruption. For SharePoint and OneDrive, files are written simultaneously to a primary and secondary datacenter region. Multiple types of checksums are stored in metadata in a separate location than the corresponding files and are used to ensure data integrity through all stages of the data lifecycle.

Because of the built-in data resiliency checks and processes, Microsoft maintains backups only of Microsoft 365 information system documentation (including security-related documentation), using built-in replication in SharePoint Online and our internal

code repository tool, Source Depot. System documentation is stored in SharePoint Online, and Source Depot contains system and application images. Both SharePoint Online and Source Depot use versioning and are replicated in near real time.

Resources

- Exchange Online Database Availability Groups
- How does Exchange Online protect mailbox data?
- SharePoint & OneDrive storage resilience
- SharePoint & OneDrive integrity checks

Data monitoring and self-healing in Microsoft 365

Article • 03/02/2023

Given the scale of Microsoft 365, it would be impossible to keep customer data resilient and safe from malware without built-in monitoring that is comprehensive, alerting that is intelligent, and self-healing that is fast and reliable. Monitoring a set of services at the scale of Microsoft 365 is very challenging. New mindsets and methodologies needed to be introduced, and whole new sets of technology needed to be created to operate and manage the service in a connected global environment. We have moved away from the traditional monitoring approach of data collection and filtering to create alerts to an approach that is based on data analysis; taking signals and building confidence in that data and then using automation to recover or resolve the issue. This approach helps take humans out of the recovery equation, which in turn makes operations less expensive, faster, and less error prone.

Fundamental to Microsoft 365 monitoring is a collection of technologies that comprise our Data Insights Engine, which is built on Azure, SQL Azure, and open-source streaming database technology. It is designed to collect and aggregate data and reach conclusions. Currently, it processes more than 500 million events per hour from more than 100,000 servers (~15 TB per day) scattered across dozens of datacenters in many regions, and these numbers are growing.

Microsoft 365 uses *outside-in monitoring*, which involves creating synthetic transactions to test everything that is important. For example, in Exchange Online each scenario is testing every database worldwide every five minutes in a scattered fashion, providing near continuous coverage of everything that lives in the system. From multiple locations, 250 million test transactions per day are performed to create a robust baseline or heartbeat for the service.

Microsoft 365 also uses the concept of *Red Alert*, which shrinks down all the monitoring signals from all of the machines in our datacenters to something manageable by a human being. The concept is quite simple: If something is happening across multiple signals, there must be something going on. It is not about building confidence in one signal, it is about having reasonable fidelity for each signal so that you get greater accuracy. This monitoring system is so powerful that we do not have 24x7 staff watching our monitors; all we have is the machinery that wakes up if it detects a problem, in which case it will page the appropriate on-call personnel, or more often as is the case, it will just go ahead and solve the problem. Once we start collecting signals and building red alerts off them, we can start triangulating across all our service partitions.

Based on the combination of the failure alert and the Red Alerts, this alert indicates exactly which components could be having a problem, and that the system is going to try to correct the problem by itself by restarting a mailbox server.

In addition to self-healing capabilities such as single page restore, Exchange Online includes several features that take an approach to monitoring and self-healing which focuses on preserving the end-user experience. These features include *Managed Availability*, which provides built-in monitoring and recovery actions, and AutoReseed, which automatically restores database redundancy after a disk failure.

Managed Availability

Managed availability provides a native health checking and recovery solution that monitors and protects the end user's experience through recovery-oriented actions. Managed availability is the integration of built-in monitoring and recovery actions with the Exchange high availability platform. It's designed to detect and recover from problems as soon as they occur and are discovered by the system. Unlike previous external monitoring solutions and techniques for Exchange, managed availability doesn't try to identify or communicate the root cause of an issue. Instead, it's focused on recovery aspects that address three key areas of the end-user experience:

- **Availability** - Can users access the service?
- **Latency** - How is the experience for users?
- **Errors** - Are users able to accomplish what they want?

Managed availability is an internal feature that runs on every Microsoft 365 server running Exchange Online. It polls and analyzes hundreds of health metrics every second. If something is found to be wrong, most of the time it is fixed automatically. But there will always be issues that managed availability will not be able to fix on its own. In those cases, managed availability will escalate the issue to a Microsoft 365 support team by means of event logging.

AutoReseed

Exchange Online servers are deployed in a configuration that stores multiple databases and their log streams on the same non-RAID disk. This configuration is often referred to as *just a bunch of disks* (JBOD) because no storage redundancy mechanisms, such as RAID, are being used to duplicate the data on the disk. When a disk fails in a JBOD environment, the data on that disk is lost.

Given the size of Exchange Online and the fact that deployed within it are millions of disk drives, disk drive failures are a regular occurrence in Exchange Online. In fact, more than 100 fail every day. When a disk fails in an on-premises enterprise deployment, an administrator must manually replace the failed disk and restore the affected data. In a cloud deployment the size of Microsoft 365, having operators (cloud administrators) manually replacing disks is neither practical nor economically feasible.

Automatic Reseed, or *AutoReseed*, is a feature that is the replacement for what is normally operator-driven action in response to a disk failure, database corruption event, or other issue that necessitates a reseeding of a database copy. AutoReseed is designed to automatically restore database redundancy after a disk failure by using spare disks that have been provisioned on the system. If a disk fails, the database copies stored on that disk are automatically reseeded to a preconfigured spare disk on the server, thereby restoring redundancy.

Data resiliency in Microsoft 365

Article • 03/02/2023

Given the complex nature of cloud computing, Microsoft is mindful that it's not a case of if things will go wrong, but rather when. We design our cloud services to maximize reliability and minimize the negative effects on customers when things do go wrong. We have moved beyond the traditional strategy of relying on complex physical infrastructure, and we have built redundancy directly into our cloud services. We use a combination of less complex physical infrastructure and more intelligent software that builds data resiliency into our services and delivers high availability to our customers.

Resiliency and recoverability are built-in

Building in resiliency and recovery starts with the assumption that the underlying infrastructure and processes will fail at some point: hardware (infrastructure) will fail, humans will make mistakes, and software will have bugs. While it would be incorrect to say that software developers were not thinking about these things before the cloud, how these issues were handled in a typical IT implementation was different before the cloud:

- First, hardware and infrastructure protections were significant. This structure meant having datacenters with 99.99% reliability required significant power and network redundancy, and servers were implemented with hardware-based clustering, dual power supplies, dual network interfaces, and the like.
- Second, process was paramount. Operations teams maintained rigorous procedures, change windows were employed, and there was often significant project management overhead.
- Third, deployment took place at a glacial pace. Deploying code without owning the source meant waiting for patch releases, and major version releases involved hardware replacement and significant capital outlay. Moreover, the only way to correct a problem was to roll back. Thus, most IT organizations would deploy only major releases to avoid the work to keep up to date.
- Finally, the scale of deployed systems and the level of their interconnectedness was historically much smaller than it is now.

Today, customers expect continuous innovation from Microsoft without compromising quality, and this is one of the reasons why Microsoft's services and software are built with resiliency and recoverability in mind.

Microsoft 365 data resiliency principles

Resiliency refers to the ability of a cloud-based service to withstand certain types of failures and yet remain fully functional from the customers' perspective. Data resiliency means that no matter what failures occur within Microsoft 365, critical customer data remains intact and unaffected. To that end, Microsoft 365 services have been designed around five specific resiliency principles:

- There is critical and non-critical data. Non-critical data (for example, whether a message was read) can be dropped in rare failure scenarios. Critical data (for example, customer data such as email messages) should be protected at extreme cost. As a design goal, delivered mail messages are always critical, and things like whether a message has been read is non-critical.
- Copies of customer data must be separated into different fault zones or as many fault domains as possible (for example, datacenters, accessible by single credentials (process, server, or operator)) to provide failure isolation.
- Critical customer data must be monitored for failing any part of Atomicity, Consistency, Isolation, Durability (ACID).
- Customer data must be protected from corruption. It must be actively scanned or monitored, repairable, and recoverable.
- Most data loss results from customer actions, so allow customers to recover on their own using a GUI that enables them to restore accidentally deleted items.

Through the building of our cloud services to these principles, coupled with robust testing and validation, Microsoft 365 is able to meet and exceed the requirements of customers while ensuring a platform for continuous innovation and improvement.

Related articles

- Dealing with Data Corruption
- Malware and Ransomware Protection
- Monitoring and Self-Healing
- Exchange Data Resiliency

Service resource limits

Article • 10/20/2023

Resource limits are enforced using quotas (limits) and throttling. Microsoft Entra ID and the individual Microsoft 365 services use both. Limits are service-specific and change over time as new capabilities are added. For details on the current limits for the various services, see the following topics:

- Microsoft Entra service limits and restrictions
- Exchange Online Limits
- SharePoint Online software boundaries and limits
- Skype for Business Limits
- File Size Limits in Sway

In addition to these limits, several throttling mechanisms are used throughout Microsoft Entra ID and Microsoft 365. Throttling within the service is especially important, given that network resources in Microsoft's datacenters are optimized for the broad set of customers that use the services. Throttling mechanisms include:

- Microsoft Entra ID and Microsoft 365 feature user-level throttling, which limit the number of transactions or concurrent calls (by script or code) that can be performed by a single user.
- A default PowerShell throttling policy is assigned to each tenant at tenant creation. These settings affect other items, such as the maximum number of simultaneous PowerShell sessions that can be opened by a single administrator.
- Each Exchange Online customer has a default Exchange Web Services (EWS) policy that is tuned for EWS client operations, and throttling that applies to all Outlook clients.

Service health and continuity

Article • 10/25/2023

Microsoft admins can view the status of services and find out when maintenance is scheduled. Service health information is available at any time by signing in. If you're using Office 365 operated by 21Vianet, some of the information below might not apply. Instead, see the **21Vianet service level agreement**.

View status of services: The Service health section shows the current status of the service and details about service disruptions and outages. Planned maintenance information is available on the Message Center. For more information, see View the status of your services.

Service incidents: A service incident is an event that affects the delivery of a service. Service incidents may be caused by hardware or software failure in the Microsoft data center, a faulty network connection due to a change made by Microsoft, or a major data center challenge such as fire, flood, or regional catastrophe. Interruptions caused by third party service providers, or changes made within customer managed environment, aren't considered service incidents. Most service incidents can be addressed using Microsoft technology and process solutions and are resolved within a short time. However, some service incidents are more serious and can lead to longer term outages.

Service notifications: There are two types of notifications about times when services may not be available: **Planned maintenance events** and **Unplanned downtime.**

Planned maintenance events: Planned maintenance is regular Microsoft-initiated service updates to the infrastructure and software applications. Planned maintenance notifications inform customers about service work that might affect the functionality of a Microsoft service. Customers are notified no later than five days in advance of all planned maintenance through Message center on the Microsoft 365 admin center. Microsoft typically plans maintenance for times when service usage is historically at its lowest based on regional time zones.

Unplanned downtime: Unplanned service incidents occur when one of the services is unavailable or unresponsive due to a failure within the Microsoft managed environment. Customers are notified of known service incidents through Service health on the Microsoft 365 admin center.

Recent worldwide uptimes: Moving to a cloud service shouldn't mean losing the ability to know what's going on. With Microsoft 365, it doesn't. We aim to be transparent in our operations so you can monitor the state of your service, track issues, and have a historical view of availability. The following tables show recent worldwide uptime data.

> **Note**
>
> This data does not apply to U.S. Government DoD, and GCC High.

Year	Q1	Q2	Q3	Q4
2023	99.98%	99.99%	99.99%	
2022	99.98%	99.98%	99.99%	99.99%
2021	99.97%	99.98%	99.99%	99.98%
2020	99.98%	99.99%	99.97%	99.97%
2019	99.97%	99.97%	99.98%	99.98%

Notification policy: When a service incident occurs, Microsoft recognizes that timely, targeted, and accurate communications are critical for customers. Microsoft notifies administrators by communicating directly to impacted customers via Service health on the Microsoft 365 admin center. Service incident updates are provided on an hourly cadence or, if a different cadence is required, it will be stated in the SHD communication posting.

Service health communication channels -- Admin App: The Admin App for organization administrators gives you the ability to connect with your organization's Microsoft service status on the go. Microsoft administrators will have the ability to view service health information and maintenance status updates from their mobile devices. For more information, visit the Admin App FAQ.

Microsoft 365 Management Pack for Microsoft System Center Operations Manager: Microsoft System Center Operations Manager (SCOM) is an integrated management platform that helps you manage data center, client devices, and hybrid cloud IT environments. Microsoft administrators who use SCOM have the option to import the Microsoft 365 Management Pack, which lets them view all service communications within Operations Manager in System Center. Using this tool gives you access to the status of your subscribed services, active and resolved service incidents, and Message center communications. For more information, get the Microsoft System Center Management Pack for Microsoft 365 in the Microsoft Download Center.

Microsoft 365 Service Communications API in Graph: The Microsoft 365 Service Communications API lets you access service communications the way you want. With this API, you can create or connect your tools to service communications, potentially simplifying how you monitor your environment. The Service Communications API lets

you monitor the following items your environment: Real-time service health and Message Center communications. For more information, see the Microsoft 365 Service Communications API reference.

Post-incident reviews: Microsoft's commitment to continuous improvement involves analysis of unplanned customer-impacting service incidents to minimize future recurrence. Unplanned service incidents are defined as multi-tenant service disruptions that impact service usage as defined by our service-level agreements (SLAs), and have been declared as such through Service health on Microsoft 365 admin center. For unplanned customer-impacting service incidents in which there was broad and noticeable impact across a large number of organizations, a preliminary post-incident review (PIR) will be delivered via your Service health within 48 hours of incident resolution, followed by a final PIR within five business days.

PIR report: The detailed PIR report includes: User experience and customer impact, Incident start and end date/time, Detailed timeline of impact and resolution measures, and Root cause analysis and actions being taken for continuous improvement. For all other service incidents, the Service health page on Microsoft 365 admin center will provide an incident closure summary including a final summary of the event, root cause, start and end times, and information detailing next steps. For this category of service incident, a PIR won't be generated.

Service continuity: Microsoft offerings are delivered by highly resilient systems that help to maintain peak service performance. Service continuity provisions are part of the system design. These provisions enable Microsoft to recover quickly from unexpected events such as hardware or application failure, data corruption, or other incidents that affect users. These service continuity solutions also apply during catastrophic outages (for example, natural disasters or an incident within a Microsoft data center that renders the entire data center inoperable).

Outage recovery: After recovery from catastrophic outages, there may be a period of time before full data center redundancy is restored for the service. For example, if Data Center 1 fails, services are restored by resources in Data Center 2. However, there may be a period of time until services in Data Center 2 have service continuity support either by restored resources in Data Center 1, or new resources in Data Center 3. The Microsoft Service Level Agreement (SLA) applies during this time. Office 365 operated by 21Vianet has a different SLA. For more information, see the 21Vianet site.

Ensuring data availability: Microsoft ensures that customer data is available whenever it's needed through the following features: **Data storage and redundancy, Data monitoring,** and **Completing preventative maintenance.**

Data storage and redundancy: Customer data is stored in a redundant environment with robust data protection capabilities to enable availability, business continuity, and rapid recovery. Multiple levels of data redundancy are implemented, ranging from redundant disks to guard against local disk failure to continuous, full data replication to a geographically diverse data center.

Data monitoring: Microsoft services maintain high levels of performance by monitoring: Databases, Blocked processes, Packet loss, Queued processes, and Query latency.

Completing preventative maintenance: Preventative maintenance includes database consistency checks, periodic data compression, and error log reviews.

Reports: Administrators of Microsoft can view reports showing how your organization is using Microsoft services. You can use these reports to identify issues, filter data, and download data to Microsoft Excel. You can also create your own reports using the Microsoft 365 reporting web services. Exchange Online and Exchange Online Protection (EOP) administrators can use mail protection reports to view data about malware, spam, and rule detections. For more information, see View and download reports about service usage. For Office 365 operated by 21Vianet, see View and download reports for Office 365 operated by 21Vianet.

Built-in service resiliency in Microsoft 365

Article • 03/02/2023

As your cloud collaboration provider, Microsoft recognizes the need to continuously earn your trust by providing solutions that function consistently and that your users love. When any given service is unavailable, it is called downtime. The definition of downtime varies for each Microsoft 365 service, but they commonly focus on any period of time when users are unable to use the essential functionality of the service. For example, here's the definition of downtime for SharePoint Online taken from the Microsoft 365 service level agreement:

"**SharePoint Online Downtime**: Any period of time when users are unable to read or write any portion of a SharePoint Online site collection for which they have appropriate permissions."

You can find the downtime definitions for each service in the Service Level Agreements.

To minimize downtime, either planned or unexpected, Microsoft 365 services are designed and operated to be highly available and resilient to failure by focusing on four areas:

Active/Active design

In Microsoft 365, we are driving towards having all services architected and operated in an active/active design that increases resiliency. This design means that there are always multiple instances of a service running that can respond to user requests and that they are hosted in geographically dispersed datacenters. All user traffic comes in through the Microsoft Front Door service and is automatically routed to the optimally located instance of the service and around any service failures to prevent or reduce impact to our customers.

Reduce incident scope

The scope of a service incident is measured by how severe it is, how long it lasts and how many customers are impacted. We strive to limit the scope of all incidents by:

- having multiple instances of each service partitioned off from each other

- deploying updates in a controlled, graduated fashion using rings of validation so that any issues that might arise from the update can be detected and mitigated early in the deployment process. This design allows for regression of the update if needed and first occurs in a small group inside Microsoft (inner ring) before it is deployed for larger groups like all 140,000 Microsoft employees (ring 2), then for early adopter rings (ring 3) and ultimately for all customers globally (ring 4).
- driving improvements in monitoring through automation. Microsoft 365 is a large service, and the SLA target uptime is high. At the very beginning of a service incident, if humans had to be involved in detection and response, we couldn't respond fast enough to meet SLAs. Automation is the key to fast and effective service incident detection and response. The sooner we know about something, the faster it can be fixed.

Along with the active/active capabilities built into Microsoft 365 service architecture, these efforts mitigate the severity, duration, and number of impacted customers during a service incident.

Fault isolation

Just as the services are designed and operated in an active/active fashion and are partitioned off from each other to prevent a failure in one from affecting another, the code base of the service is developed using similar partitioning principles called fault isolation. Fault isolation measures are incremental protections made within the code base itself. These measures help prevent an issue in one area from cascading into other areas of operation. Fault isolation measures are applied at multiple stages of the development and delivery of a service, including code development, service deployment, load balancing, and database replication.

The Microsoft Security Development Lifecycle (SDL) further promotes resiliency and consists of a set of practices that support security and compliance requirements. SDL guides our developers in building resilient, secure, compliant services. Key elements of SDL include code reviews, threat modeling, penetration testing, and standardized incident response processes across the Microsoft cloud.

Microsoft 365 services are highly interconnected, but the systems and technology behind them are engineered in a way that limits the impact of one service incident from spilling over to other services. For example, an issue affecting Exchange Online will not impact core functionality in Teams, or an issue with search functionality in SharePoint Online won't affect users' ability to upload or download files.

Continuous service improvement

When we experience an incident, we take it seriously. After all, our redundant cloud architecture and rigorous internal processes aim to keep our services accessible. During an incident, our monitoring rapidly detects the affected services and, if your tenant is affected, you'll be notified immediately through various channels. Simultaneously, engineers follow well-defined processes to triage the issue and take the necessary steps to restore normal operation as quickly as possible. Once the service is functioning normally again, we hold post incident reviews as part of the cycle of continuous service improvement. During the post incident review, we identify the root causes of the incident and what was required to fix the issues. Then we take what was learned from the situation and apply it to the design and operations of all of our suite of offerings. With this knowledge, we can prevent the same root cause from impacting other services and additional customers.

Enterprise business continuity management customer and cloud partner responsibilities

Article • 03/02/2023

Getting Microsoft 365 cloud services to your users is a partnership between your organization and Microsoft. Microsoft provides the services and you are responsible for connecting your client endpoints, managing identity and access and how those services are used. There are shared responsibilities, such as the identity and directory infrastructure as well. This article covers some of the critical items you need to be mindful of to keep your business functioning in the event of a service incident and it helps set expectations as to what Microsoft will do during a service incident.

Transparency during service incidents

As a trusted partner, Microsoft builds highly resilient cloud services and follows structured procedures to resolve service incidents when they happen. When a service incident occurs, Microsoft recognizes that **timely**, **targeted**, and **highly available** communications are critical for customers.

Timely

Microsoft notifies Microsoft 365 administrators by updating the tenant-specific Service Health Dashboard (SHD) in the Microsoft 365 Admin Portal. Service incident updates are normally provided on an hourly cadence. If a different cadence is needed, we'll keep you informed of the change in the SHD communication postings.

Targeted

In most cases, when our monitoring systems detect an issue, we can identify the affected customer base, from a single customer up to region or beyond and direct the necessary communications to those customers. This helps you know what you need to know for your business and not be distracted by noise notifications that don't impact you. For example, if a specific mailbox database is impacted, we're able to identify exactly which customers have users on the affected infrastructure and scope our communications to them. If the scope of impact of the incident is unclear, we expand

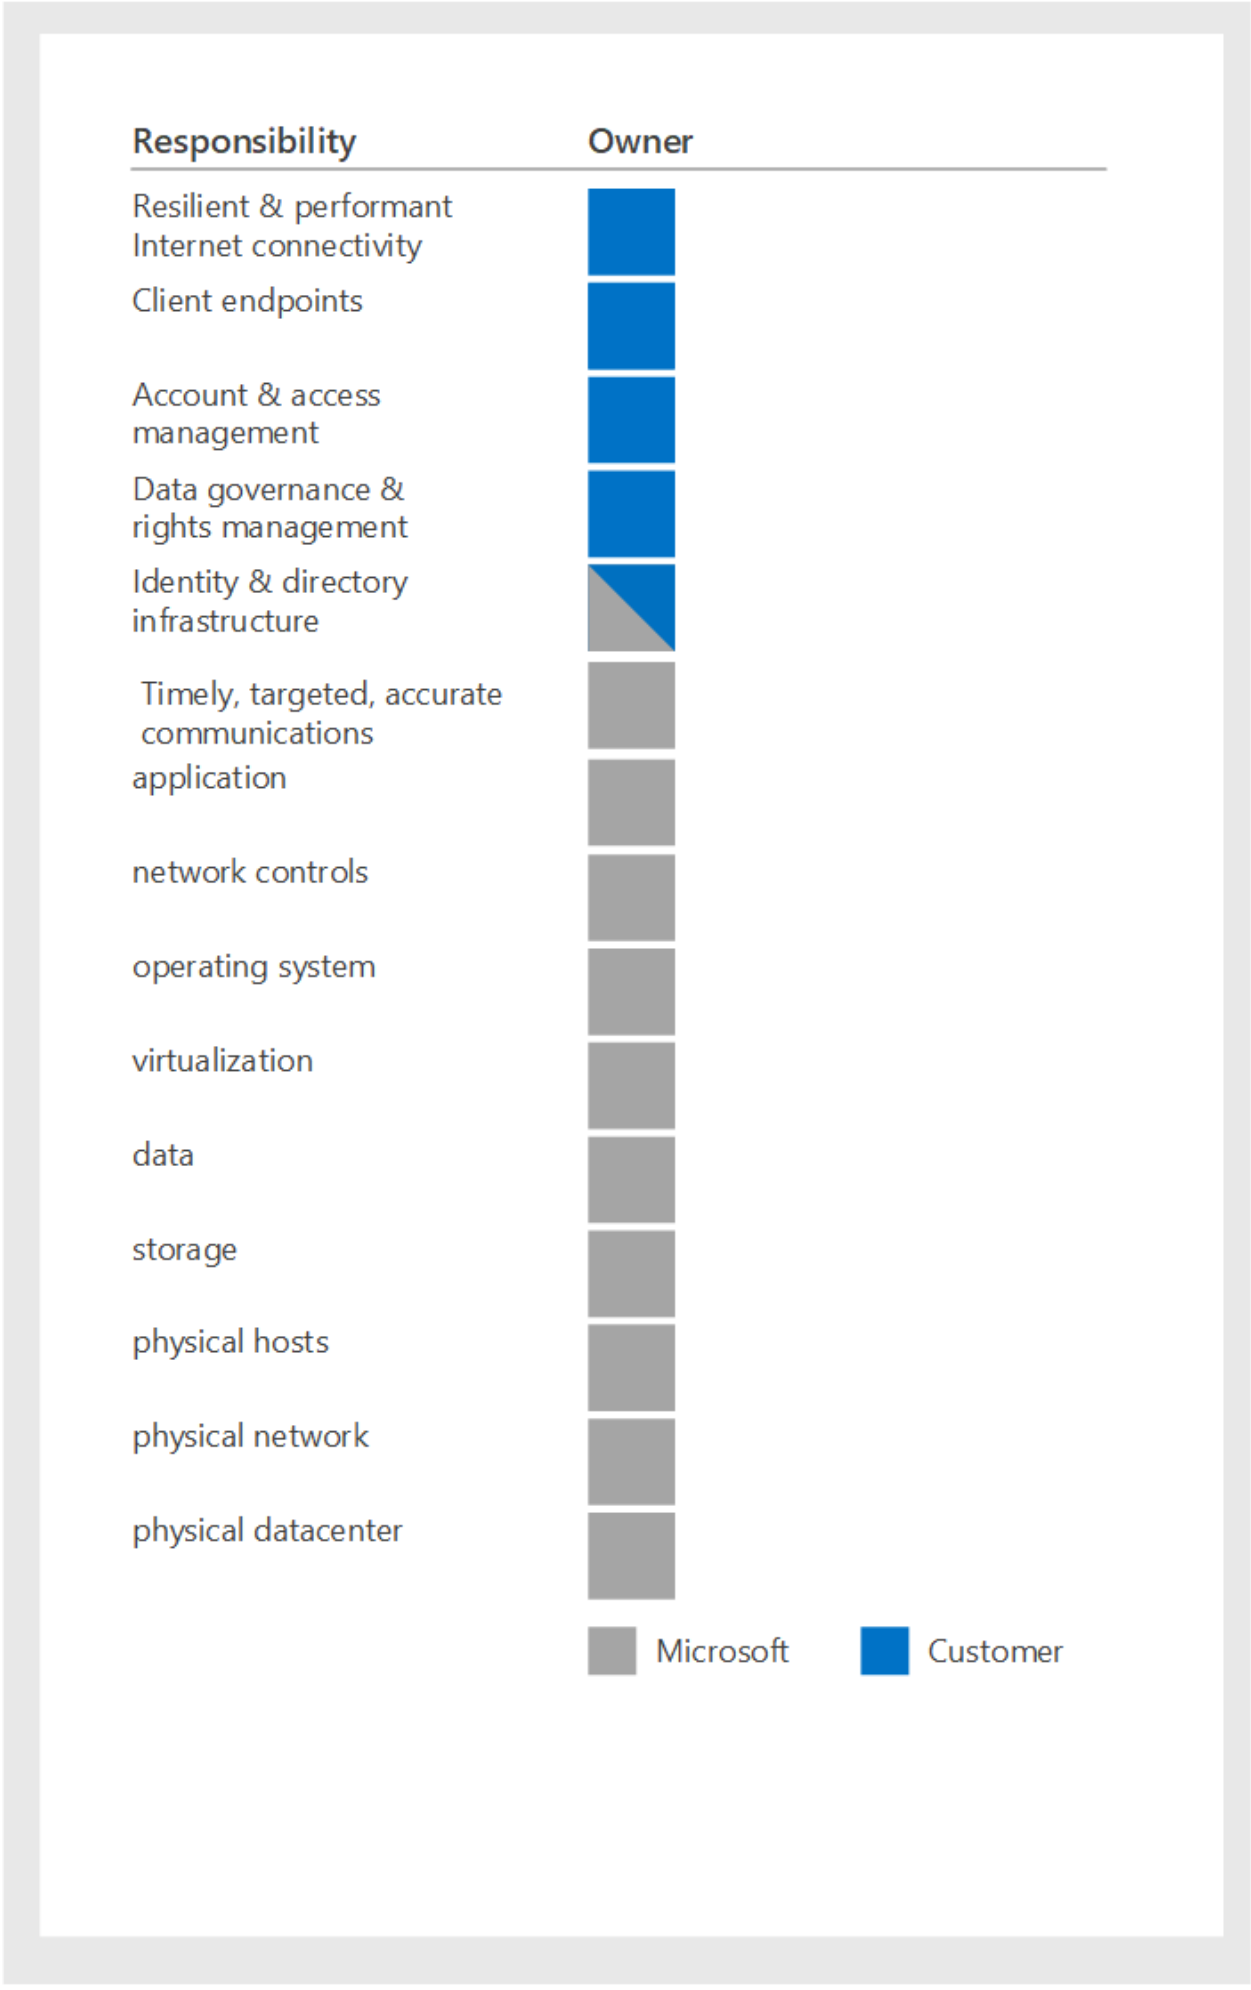
our communications out to the widest group of customers who are potentially impacted.

Highly available

Microsoft maintains multiple channels for service status communications that customers can use.

- In the event the Admin center or the Service Health Dashboard within the Admin center are unavailable, you can monitor the service status using our backup site.
- We maintain a Twitter account @MSFT365Status where we will respond to reports of impact and post updates on SHD impacting events.
- The Admin App for Microsoft 365 tenant administrators gives you the ability to connect with your organization's Microsoft 365 service status on the go. Tenant administrators will have the ability to view service health information and maintenance status updates from their mobile devices. For more information, visit the Admin App FAQ.
- The Microsoft 365 Service Communications API enables you to access service communications so you can more easily monitor your environment. You can connect to the API, receive real-time service health data, and publish the information on an internal dashboard to inform enterprise users of incidents. Distributing the information internally can decrease your helpdesk traffic during an outage.
- For major incidents, Microsoft publishes Post Incident Reviews (PIR) to the SHD within the Admin center. PIRs contain key incident information to help you understand the nature of the outage. It generally contains the following sections:
 - User impact
 - Scope of impact
 - Incident start-end date and time
 - Root cause
 - Actions taken
 - Next steps
- Ancillary communications are available in the Microsoft 365 Message Center, such as notices of upcoming changes, new features, or planned maintenance.
- For more information, see the Service Health and Continuity guide to learn more about the different communication channels and how to monitor service health.

Providing access to Microsoft 365 online services is a partnership between your organization and Microsoft. The following chart summarizes the balance of responsibility for both Microsoft and the customer during a service incident and during regular operations.


Responsibility
Owner
Resilient & performant Internet connectivity
Client endpoints
Account & access management
Data governance & rights management
Identity & directory infrastructure
Timely, targeted, accurate communications
application
network controls
operating system
virtualization
data
storage
physical hosts
physical network
physical datacenter
Microsoft
Customer

Your environment - service continuity

When thinking about your continuity plan, be mindful of events that may impact your organization and its overall ability to communicate. At a high level, there are three factors that could impact your business.

People

Consider events that would cause impact to your workforce like a natural disaster or a pandemic. This is often overlooked, due to the unlikely nature of a broad-scale impact if your workforce is widely distributed. But, if a large percentage of the workforce were offline, would your business continue to operate? How do you mitigate that?

Location

Many organizations require employees to be in specific physical or network locations in order to connect to enterprise systems and cloud services.
Microsoft publishes network connectivity principles that guide enterprises through best practices for setting up network connectivity to cloud resources. Examples of optimization include implementation of split tunnel VPNs to allow connections directly from a user's network rather than over a VPN tunnel. While these connectivity principles are important for maintaining low-latency connections, service resiliency requires alternative methods of connecting to corporate resources for general collaboration.

Systems

Many collaboration solutions are dependent on systems, such as the company-wide area network (WAN). When those systems are not available, how would your organization respond? This graphic represents issues that may impact more than one area. The accompanying table provides examples to consider



	LAN failure	Pandemic	Severe weather event	Security Breach
Systems	Yes	Yes	?	Yes
People	No	Yes	Yes	No
Location	No	Yes	Yes	No

Your continuity plans should consider each of these areas. For example: If you require users to be on the corporate network and there is a snowstorm, how do those users

gain access to key resources? If the snow prevents travel into the office and service engineers are required to connect to the corporate network, is there a policy mandating they have their corporate laptops in their possession at home?

Developing your business continuity plan

Article • 03/02/2023

This topic provides guidance on developing a business continuity plan that takes Microsoft 365 dependencies into account. Here we recommend methods for analyzing your business functions and identifying the ones that depend on Microsoft 365 services. You'll perform this analysis with the anticipation that there will be service failures and that you have to prepare for those possibilities.

Broadly speaking, business continuity planning involves four aspects, assessment, planning, capability validation, and communication and coordination.

Assessment

First you must identify the business functions in your org and the services and processes that support them. This includes completing a business impact analysis, where each business function is ranked according to how critical it is and you identify the processes and services that each one depends on. Here's a sample table you can refer to help you get started with your own assessment.

Sample Business Impact Assessment (BIA)

This is a BIA document for `name of the service, system, process, or function`

BIA fields	Description
BIA type	`is it a business process or technology, service or system?`
BIA name	`name of the service/system/function/process`
Service description	`give a full description of the service, process, or function`
Enterprise function	`some examples: customer services; legal; marketing; risk management, security, sales, information technology, production, manufacturing`
Fiscal year	`the current fiscal year, re-evaluate these on a regular basis`

BIA fields	Description
Criticality	`develop your own classifications, but here are some examples: mission critical, important, deferrable`
Business unit	`name of the business unit that owns this business function`
Process (service, feature)	`the name of the process, service, or feature`
Business group senior leader	`the name and contact information of the senior leader of the business group that owns this business process`
Does the technology have an established **internal** SLA or OLA?	`please explain in as much detail as possible`
Does the technology have an established **external** SLA or OLA?	`please explain in as much detail as possible`
Does the technology have a known executive mandate driving a specific process SLA? If yes, explain in detail.	`details here`
Will the loss or compromise of the data associated with this service trigger a major event? If yes, explain in detail.	`details here`
Does the service have a workaround or alternative in place for some or all of its key functions and features? If yes, explain in detail.	`details here`
Does the service process, store, or transmit customer data, such as personally identifiable information (PII)? If yes, explain in detail.	`details here`
BIA status	`develop your own status classification, here are some examples: planned, started, in-progress, complete, on-hold, expired`
Completion date	`the date this BIA was completed`
BIA facilitator	`name of the person or group who is responsible for developing and maintaining this BIA`
BIA approval	`name of the person or group who is the executive sponsor of this BIA and who has responsibility for approving it.`

BIA fields	Description
Contributors	`optional list of the people who helped develop this BIA and their contact information`
BIA approval location	`indicate where the executive approval is located, or attach proof to this document`

Planning

Next, you look across business processes to see where any cascading dependency relationships exist. Based on the outcome, you prioritize and form resiliency strategies, and standard operating procedures supporting your strategies.

You can use Microsoft Service Map to help you in with this mapping. Microsoft Service Map automatically discovers application components on Windows and Linux systems and maps all TCP dependencies, identifies connections, and remote third-party systems that the app depends on. It also maps dependencies to areas of your network that are traditionally dark, such as Active Directory.

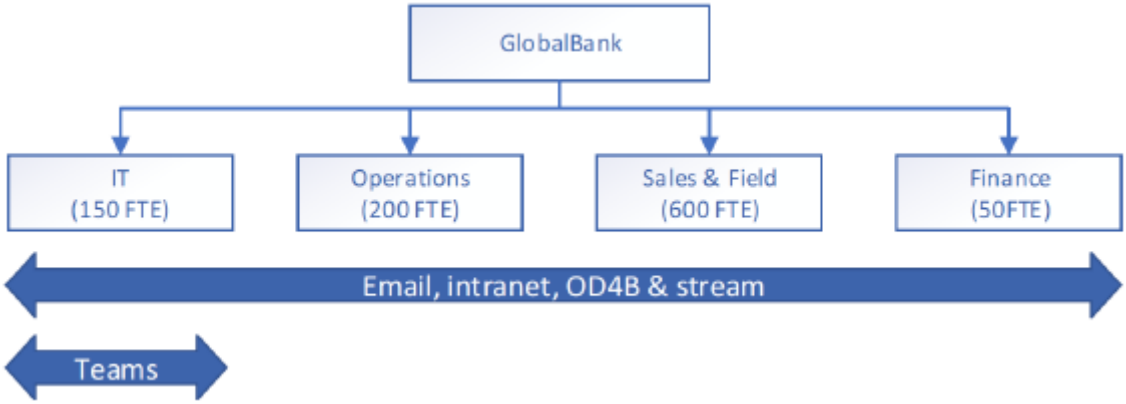
Here's a sample dependency analysis (DA) you can start from. In your dependency analysis (DA), you will identify and examine the process dependencies. Make sure you include people, suppliers, customers, partnerships, and facilities. The data from this analysis will be used to identify gaps between the recovery requirements of a process and the recovery capabilities of supporting dependencies.

Field	Description
Process type	
Facilitator	
Completed by	
Completed date	
Contributors	

Capability validation

Once you have inventoried your business processes and mapped out relationships to other process and technologies, you need to build validation scenarios for all the processes. Basically, figure out how you are going to validate your business process continuity plans. You'll probably find that some are more important that others and

you'll want to prioritize those. Don't forget that regularly training for employees on incident response and continuity measures is important once the plan is established. Post incident reviews should be used to enhance your resiliency strategies by incorporating learnings from each validation or test.



Incident coordination and communication

During a service incident, normal communications channels may be impacted or degraded, so you should pre-arrange alternatives to help your organization stay connected during an incident. It is critical that the communication channels be established, vetted for security and compliance, and users trained on their use prior to a disruption. Failing from a known state to another known state is far preferable to users coming up with ad-hoc, unknown solutions in the middle of a crisis.

At Microsoft, each service team has established internal alternative communication channels to help us coordinate when our normal communications channels aren't available. These include backup telephony and audio-conferencing solutions, Viva Engage groups, Teams groups, internal Service Health Dashboards, and internal Incident Management software.

During your Business Impact Analysis and Dependency Analysis, you will be mapping critical processes and the technologies or services they depend on. Pay special attention to communication during this phase of planning and think of alternatives. Here are some examples.

- If email is your primary method of keeping your users and stakeholders informed, and your email service is degraded or unavailable, you can use another service such as Microsoft Teams, Viva Engage, or another 3rd-party service as a backup. The key is to establish these beforehand and train your users on where to go. A Viva Engage thread isn't going to be useful if no one knows it exists or if no one has it bookmarked.

- If your internal Incident Management processes rely on voice communications to coordinate your responses, establish an alternative telephony solution for use during a crisis. This solution doesn't need to have full parity with your primary service but should provide the minimum level of collaboration to coordinate your Business Continuity and Incident Management teams. Additionally, asking users to publish their mobile phone numbers in your Global Address List can provide an additional layer of backup communication in extreme cases.
- You may want to create a custom service health dashboard, or other such site, which can provide status updates during an incident. Training users where to go for information beforehand will help reduce unnecessary calls to help desk and instill confidence in your user base that the situation is being handled quickly and efficiently. Use the O365 Service Communications API to tie this information into Microsoft 365 for an even greater level of visibility.
- It is critical that the location of your Business Continuity Plans and Standard Operating Procedures is well known. We recommend maintaining online and offline copies of critical documentation, such as with SharePoint Online or OneDrive for Business configured for automatic sync to local devices. For Service/Network Operations Centers and other similar teams that will be critical for recovery, you may also want to keep hard copies available to be used in case event of an emergency.

Know your external points of integration

Regardless of business model, every company has points of integration with their customers, partners, and vendors. The business value supply chain is built on integration with external entities. Improving business continuity for service disruptions requires consideration and protection of each point of integration.
As you analyze your supply chain, external communications should be considered in the same way internal communications are analyzed. Do your customers rely on your Exchange Online servers as the only method of contacting you? Have you established and made your suppliers aware of alternative communication methods, in the event uptime is impacted? Here's a sample table that suggests how to organize your thinking.

External entity name	Impacting incident scenario	Microsoft 365 services integrated	Alternatives
`vendor name`	Mail flow	Exchange Online is the only means of communication with Contoso	set up external Microsoft Teams channels or a third-party collaboration software

External entity name	Impacting incident scenario	Microsoft 365 services integrated	Alternatives
`service supplier name`	Chat	Microsoft Teams	Third-party instant messaging
`partner name`	Voice	Microsoft Teams	Mobile or public pstn
`supplier name`	File sharing	Externally shared SharePoint sites and OneDrive	Third-party file sharing

Train for business continuity in your organization

Article • 03/02/2023

Ultimately, your business processes are realized by interactions between your employees, partners, and customers. You can build a terrific business continuity plan, but if you don't train your users, administrators and partners on it, it won't be nearly as successful as it would be if you had. Train your users when they are hired and run refresher training annually. Your technical staff should regularly review and make the necessary updates.

Be sure to include these key topics in your service resiliency training:

- the locations of critical business continuity resources such as call tree outlines and emergency procedures
- clear definitions of all participant roles and responsibilities during an incident
- Incident policy guidelines and standard operating procedures
- Available alternative forms of communication, such as emergency phone bridges, social media pages, or mobile phones
- Alternative means of connectivity, including the use of Microsoft 365 mobile or web applications in the event rich clients are unavailable

Plan the exercise and exercise the plan

You should regularly practice your continuity plans and update them based on what was learned from the practice. If you don't test them, your continuity plans quickly become outdated and most tabletop exercises are insufficient preparation for addressing the issues you encounter in real world test. A standard practice is to exercise each Business Continuity Plan (BCP) within 12 months (at a minimum) of the last validation. Microsoft recommends integrating continuity exercises into your routine operations more frequently, which may involve taking systems offline or temporarily disabling primary methods of communication.

When considering test scenarios, you will need to balance the potential impact to your business of injecting faults into your production environment against the benefit of highly polished business continuity processes and skillsets. During each exercise, pay special attention to any assumptions that turned out to be incorrect, and identify any gaps in the processes you've established. A failed Disaster Recovery test exercise isn't necessarily a bad thing, as it will allow you to identify problems in the plan and correct

them before a real disaster strikes. Make sure you include these points in your exercise planning.

1. Ensure that all BCP documentation is updated and available in secure, redundant repositories.
2. Train all users and stakeholders regarding the processes and procedures you've established for the drill.
3. Establish the goals of the exercise, define success criteria, and choose whether you will notify users prior to the test or not, as this will reveal whether prior training has been effective.
4. Record a detailed timeline during the test to track all actions taken by the Incident Management team.
5. Test backup communication services to ensure they can fulfill the needed role.
6. Test your call-tree to ensure that critical resources can be engaged using the alternative solution.
7. Determine if the test was a pass or fail according to the success criteria established earlier. Any failed test should be redone within 3-6 months.
8. Perform a post drill review to identify any missteps, or unanticipated gaps or issues that arose during the incident.
9. Track any gaps or issues in your internal issues-management system and assign them to the proper owner for remediation.

Service incident mitigation strategies

Article • 10/20/2023

Here are some strategies and scenarios that show how to mitigate the impact of a Microsoft 365 service incident on your business process.

Service incident scenarios and potential mitigations

Microsoft 365 dependency	potential mitigations
Incident Management system relies on Exchange Online to engage On-Call Engineers and Incident Managers.	Ensure that your Incident Management system supports multichannel communications, such as parallel email, phone call, and SMS notification, and call tree hierarchies in case the primary on-call doesn't engage, the system automatically engages the backup. Also include backup contact methods in every notification, so that backup communication methods are embedded for easy reference. Alternative communication methods, such as Viva Engage, can be used for emergency collaboration if the incident management service is unavailable.
Microsoft Teams is used for storing files accessed via the client.	Teams stores files uploaded to the client in a SharePoint Online document library. Files are still accessible via SharePoint Online. Train users on file locations in SharePoint Online.
Microsoft Teams conference calling is relied upon for general communication and incident management triage.	Establish a backup conferencing solution with a 3rd-party provider.
VoIP phones are used as a secondary method of communication.	Implement non-VoIP phones capable of PSTN calling, especially for network and service operations centers during incidents. Add employee mobile phone numbers to the company directory for enabling critical personnel to be contacted over the cellular network.
OneDrive for Business is relied upon for file storage and user productivity. Files On-Demand is configured to free up space on local user drives.	OneDrive sync supplies group policies allowing admins to require specific content to be synced locally or free up space when desired. To mitigate the risk of document inaccessibility, configure this policy to sync critical documents locally. Train users to manually apply the 'always keep on this device' setting for key documents.
Communication of business disruptions to customers	Public third-party social networks can be used as an alternative means of mass communication.

Microsoft 365 dependency	potential mitigations
and suppliers rely on Exchange Online.	
Hybrid on-premises architecture, such as ADFS or Pass Through Authentication, fails causing disruption to user's ability to authenticate to cloud services.	Configure Password Hash Sync, in conjunction with your hybrid authentication services, as a secondary cloud-based authentication mechanism to avoid sign-in disruption during the outage. Refer to Create a Resilient Access Control Management Strategy with Microsoft Entra ID for further information on building resilient authentication and access control architectures.

Leveraging mobile app access

As mobile use has proliferated, there are new means to stay connected and Microsoft 365 mobile applications can be a key part of your resiliency strategy. Because they connect to cloud services over the cellular provider network, they are not dependent on your organizations network infrastructure.

Let's use Outlook as an example. Users can Connect to their Exchange Online mailboxes over different network protocols (https or MAPI) depending on the email app being used. If there is a service incident that involves one of the protocols, say MAPI for instance which the desktop client uses, then your users can still get to their mailbox through the Outlook Mobile app or Outlook on the Web.

If you decide to allow users to connect to Microsoft 365 services via their mobile devices you can use Microsoft Intune to securely configure and manage those devices. Once the user accounts and devices are enrolled in your mobile management solution ensure that the apps have been downloaded and configured.

Strengthening operational resilience and reducing concentration risk in financial services

Article • 08/07/2023

This article builds upon our July 2023 blog post which provides practical guidance on how to strengthen operational resilience and manage concentration risk in financial services institutions (FSIs).

To ensure consistency with existing and upcoming regulations, we incorporated into our approach the main regulatory guidance for FSIs, including ongoing consultations on the topic of operational resilience.

These include:

- Financial Stability Board (FSB) consultation on enhancing third-party risk management and oversight (22 June 2023 FSB Consultation)
- The EU Digital Operational Resilience Act (DORA)
- DORA Draft Regulatory Technical Standard on ICT risk management tools (DORA ICT RMF)
- UK PRA Discussion Paper 'Operational resilience: critical third parties to the finance sector' (UK DP3/22)
- US Interagency Guidance on Third-Party Relationships
- US Treasury Report on Cloud Computing
- Canada Third Party Risk Guidelines
- Monetary Authority of Singapore's establishment of a Cloud Resilience Forum
- 2021 IOSCO Principles on Outsourcing

Our approach is consistent with the listed publications and will be updated over time. Our objective is to help customers grow and innovate in a responsible and compliant way, in line with existing and future FSI regulations.

Overview

When it comes to the use of cloud technology in FSIs, operational resilience and concentration risk are interlinked because it's addressed by the strengthening of operational resilience in various regulations. Both are concerned with a wide set of measures including the identification and monitoring of critical third-party relationships,

requirements to strengthen risk governance and management, and guidance around business continuity and exit planning.

We aim to help FSIs address and strengthen their operational resilience to help them manage concentration risk at the enterprise level in a way that is consistent with regulatory guidelines. It's important that this topic is addressed holistically and in consideration of all third-party relationships. The use of cloud services and dependencies on on-premises software products must be maintained and kept secure. In addition, subcontracting in relation to critical functions is also commonly addressed in regulatory guidance.

In the 22 June 2023 FSB Consultation, a critical service was defined as a service whose failure or disruption could significantly impair a financial institution's viability, critical operations, or its ability to meet key legal and regulatory obligations. It's primarily these services that are in focus when it comes to managing concentration risk.

Concentration must be distinguished from concentration risk. Concentration of third-party dependencies is commonplace today across many critical FSI services, and it may not be feasible to eliminate it. For instance, if you consider the broad usage of financial information networks (for example, Bloomberg or Thomson Reuters) and software vendors such as IBM and Oracle, it becomes clear that complete substitution would be challenging. One can devise an exit plan for a specific third-party service, but it becomes difficult to discontinue all use of a third-party solution for large financial institutions. Therefore, the focus should be on reducing concentration risk so that it stays within an organization's risk tolerances.

While some of these examples are already systemically important, they aren't often perceived as carrying elevated levels of concentration risk because they're distributed across multiple locations that are geographically dispersed. As a result, most failures only have a limited (nonglobal) impact. They also offer the highest levels of resilience through state-of-the-art design such as secure and resilient operations, use of automation, and implementation of zero trust. These designs have been confirmed by multiple independent third parties through various certifications processes. This collective set of measures has brought down the overall risk profile even though, at a firm level, concentration is still in place.

FSIs must focus on strengthening operational resilience, and it must be done in a way that is consistent with regulations and guidelines. It's important to note that regulations don't force FSIs to deploy hybrid or multicloud solutions but instead take a principled, risk-based, and technology neutral approach to addressing concentration risk. These risks may also exist in an on-premises environment.

Beyond operational aspects, regulations also highlight nonoperational risks in relation to outsourcing, in particular financial insolvency risks and resolution approaches. Nonoperational risks can't be managed through technological measures but instead are addressed with legal measures during the contracting phase and by invoking the exit strategy.

A 6-step approach for managing operational risk in financial services

There are several elements that must be considered before making critical choices on how to manage operational risks, which is why we introduced a six-step approach to addressing concentration risk and strengthening operational resilience:



Step 1 - Update cloud risk governance

Cloud technology in financial services regulation today is mainly governed by third-party outsourcing guidelines that apply specifically in this context, with different sets of regulation applying to the on-premises ICT environment. More recent consultations, as referenced in the introduction, **take a holistic approach when it comes to strengthening operational resilience**. For instance, draft regulations such as DORA and the DORA ICT RMF won't only apply to on-premises but will also consider the use of ICT third-party service providers as an integral part of its framework.

Additionally, as described in the June 23 FSB Consultation, **the primary focus is on critical services**. We also see this in other regulations such as DORA. The focus on

critical services is necessary because otherwise the scope of services to be assessed becomes too broad as technology services become increasingly interconnected.

Therefore, we recommend that firms revisit internal risk governance frameworks and holistically incorporate these new operational resilience measures with a focus on critical services, incorporating various guidelines on third-party or cloud outsourcing to ensure compliance.

Some appropriate questions to ask in this context include:

1. Have clear organizational risk tolerances been defined?
2. Which services are business critical?
3. What plausible to real threat scenarios may impact these?
4. What is my firm's overall risk appetite?

Risk governance frameworks must also be updated annually to ensure continued compliance in a rapidly evolving regulatory landscape.

The regulations and guidelines from the introduction align to these questions as a starting point when it comes to strengthening operational risk. To account for these requirements across jurisdictions, Microsoft has created an extensive set of Financial Services Compliance Checklists to help customers self-assess regulations in various countries when it comes to the use of cloud technology. These checklists feature regulatory mappings and point to specific information relevant when assessing the use of Microsoft cloud technology.

In addition, the Compliance Program for Microsoft Cloud has been designed to help risk and compliance functions across all three lines of defense in complying with these regulations and in addressing overall risks related to the use of cloud. The program offers both proactive and reactive features, offering a premium support channel for risk stakeholders.

Step 2 - Identify concentration

The higher the level of concentration of services with a single third-party provider, the higher the potential of adverse impact if something goes wrong, which is referred to as concentration risk. This risk is why organizations need to have a clear understanding of all dependencies between business processes, ICT platforms, software, and third-party relationships.

Mapping these dependencies is typically performed as part of the business impact analysis (BIA). Once all third-party relationships have been identified and mapped to critical use cases, it becomes possible to identify the level of concentration of critical

services with a single third-party provider. This view enables firms to identify where they must focus when managing concentration risk.

For Microsoft Cloud services, it's possible to identify potential concentrations for critical workloads by reviewing subscription and tenant IDs in the Azure portal. Firms may also view and filter Azure resource information that can help provide a detailed understanding of which services are being used within an organization. This information is exportable and can be correlated against internal BIA information to help identify critical dependencies in real-time. For Microsoft 365, firms can also generate reports in the Admin Center both on purchased licenses and on activity.

Step 3 – Assess alternatives

Once a firm understands its dependencies and how they relate to critical use cases, the next step is to address associated concentration risk. Start by identifying a short list of feasible alternatives that can be investigated further in depth. Questions you may want to address include:

- Which practical alternatives exist across on-premises, hybrid, multi-sourcing, and full cloud?
- What are the drawbacks and benefits of each?
- How do their risk profiles compare with each other? How resilient is each alternative?
- Which ones best fit the organization's risk appetite and cloud strategy?
- When planning an exit, is a full vendor exit possible and desirable? What alternatives can be considered?

Concentration risk is an aggregate term pointing to the higher impact an adverse event would have on one or more critical services. When assessing such risks, one must evaluate each of the underlying threat scenarios, which in turn leads to a nuanced view that includes both benefits and drawbacks associated with concentration of services. Threat factors to consider include data center disasters, hardware failures, network outages, cyber-attacks, faulty changes and upgrades, human errors etc. For each of these appropriate mitigating measures should be carefully considered. Firms must also consider mitigation costs, complexity, and availability of in-house skills when considering the preferential solution for addressing concentration risk.

It may be possible and, in some cases, even desirable to maintain concentration so firms can maximally strengthen resilience. **Rather than to try and remove the third-party dependency entirely, firms should focus on strengthening operational resilience** by addressing the underlying threat scenarios associated with concentration risk (for example: a regional data center disaster event). These scenarios can often be easily

addressed and with less drawbacks by (i) reducing the probability that the threat event occurs and (ii) limiting its impact by reducing concentration:

1. **Reducing probability** is achieved by strengthening resilience in the solution design. A robust set of risk management procedures can enhance operational resilience despite concentration of critical functions with a single third-party provider. Measures may include running on state-of-the-art infrastructure, running a zero-trust security model, patching systems with the latest updates, ensuring business continuity measures have been set up and proven to work, deploying modular and open-source technologies, (for example: containers) etc. Each of these contributes towards obtaining a maximally resilient environment.

2. **Limiting impact by reducing concentration at lower levels** is achieved by designing your services to operate across multiple availability zones in an active/active configuration; by ensuring sufficient redundancies and recovery mechanisms are in place (for instance backups), and by leveraging geo-redundant designs. Such configurations not only result in higher resilience and better SLAs, but also help to mitigate against threats such as the loss of a single data center or even an entire region due to their distributed nature. The impact of threats can be reduced, hereby also reducing concentration risk, in some cases even going beyond what is feasible in on-premises or hybrid scenarios.

In conclusion, **if a full cloud topology offers higher resilience compared to alternatives, concentration risk will also be effectively reduced although concentration itself is not diminished**. This outcome can be both acceptable and even desirable.

We recommend incorporating this operating model into a corporate cloud policy as this provides guidance to business and ICT teams on what the preferred topology is for an organization and which elements to consider as part of their solution design. Another reason to consider this is because strategic alliances are often made with one or more cloud providers for the delivery of ICT third-party services and approaches for managing associated risks are often managed at a higher level.

Step 4 – Design for resilience

At this stage your ICT technology, security and operations teams start to dive deep into the solution design by building the conclusions and requirements from previous into individual user cases.

ICT Teams must ensure that they configure and design applications to be secure and resilient by default. This means ensuring that solutions are reliable, secure, free from

single points of failures, and potentially leverage availability zones to establish recovery time objectives (RTOs), recovery point objectives (RPOs), and service levels (SLAs) as required by the business. Implementing backups where necessary, updating business continuity and exit plans is also part of the process.

When trying to maximize the end-to-end SLAs with Microsoft cloud technology, consider reviewing our SLAs for Microsoft Online Services. You'll find that SLAs vary by service and are dependent upon customer design choices, with higher SLAs for deployments across multiple availability zones. By choosing the right design, customers may achieve a 99.999% monthly availability SLA in the cloud.

Ensuring a strong and secure-by-default design is no easy task and there's a risk that ICT teams leave weaknesses or make errors during their implementations. This challenge is why we offer guidance on best practices when deploying on the Microsoft cloud. While it's less complicated for SaaS offerings such as Microsoft 365 and Dynamics, things can become complex when a solution is built on Azure (IaaS). Therefore, the Microsoft Azure Well-Architected Framework was created by Microsoft to provide additional guidance on how to design for reliability and security. Microsoft provides resources to aid firms in implementing high resilience scenarios, including an overview of the Azure reliability, the overview of Azure security and review the concepts of Azure Availability Zones and Regions. Other aspects to consider include establishing operational excellence and to a lesser extent optimizing cost and performance.

A lot of attention must therefore go to overall security, especially in financial services. We recommend evaluating our Security in the Microsoft Cloud Adoption Framework for Azure, which provides core principles and practical guidance on how to address today's cybersecurity threats most effectively on our cloud. It also includes links to reference architectures and security baselines for different use cases.

Finally, ICT Teams should deploy cloud resources in a way that meet all regulatory and internal policy requirements. Azure Governance helps firms implement these regulatory and internal control requirements by enforcing such policies to Azure cloud resources. In addition, firms may review our introduction to Azure hybrid and multicloud, which helps support hybrid and multicloud deployments.

Step 5 – Test business continuity plan

The process of business continuity planning and testing has been well established in regulated FSIs. The focus in this step is on evaluating the impacts that may cause disruption of the third-party outsourcing (use of Microsoft cloud) on this process. A good example of testing is addressed in Article 26 *Testing of the ICT business continuity*

plans of the DORA draft RTS on ICT risk management tools methods processes and policies.

There's an aspect of shared responsibility to this, and our Microsoft resiliency and continuity overview helps customers prepare for disaster scenarios. It offers specific guidance for each Microsoft cloud service, including how Microsoft deals with business continuity for each service and guidance on how customers can leverage Microsoft cloud services to prepare for disaster events.

Business continuity should be tested regularly by financial firms when using Microsoft Cloud as the responsibility follows the shared responsibility model. For SaaS solutions such as Microsoft 365 and for some Azure services, Microsoft is responsible for performing these tests at regular intervals. For more insights, customers may review Microsoft's quarterly Business Continuity and Disaster Recovery Plan Validation Report that Microsoft publishes on its Service Trust Portal under the Business Continuity and Disaster Recovery section for more insights on how specific Microsoft cloud services have performed.

Step 6 – Prepare exit plans

Some threat scenarios can't be managed with business continuity plans or technical resiliency measures, such as the risk of bankruptcy or resolution of the third-party provider. An exit plan has the benefit of dealing with such catastrophic scenarios and should be seen as complimentary to having tested business continuity plans.

This is why every organization should have an overall exit strategy and individual exit plans for its critical use cases as this is grounded in several regulations and future guidelines.

Section 3.7 of the June 23 FSB Consultation offers helpful guidance pointing to elements of an exit strategy and exit plans such *(i)* contractually agreeing on transition periods to minimize the risk of disruption; *(ii)* ensuring that logical and physical assets including data and applications are returned in a cost-effective and timely manner and *(iii)* having contractual provisions relating to the ownership, maintenance, preservation, and long-term availability of records as appropriate.

DORA Article 28 (8) also requires that firms develop exit strategies for ICT services supporting critical or important functions. Similarly, the UK PRA's Supervisory Statement SS2/21 on Outsourcing and third party risk management of March 2021 has an extensive description of exit planning requirements in section 10 and also introduces the concept of stressed exits. Financial firms in these jurisdictions must evaluate these guidelines and put in place adequate contractual arrangements that can support the

execution hereof, for instance by requiring mandatory transition periods during which ICT third-party service providers continue providing relevant services (example: DORA Article 30.3.f references the contractual establishment of a mandatory adequate transition period).

A technical paper written by the European Banking Federation in June 2020 describes some of the ways exit plan testing can be achieved for financial institutions in line with EBA Guidelines.

Security development and operations overview

Article • 04/04/2023

How does Microsoft implement secure development practices?

Microsoft's Security Development Lifecycle (SDL) is a security assurance process focused on developing and operating secure software. The SDL provides detailed, measurable security requirements for developers and engineers at Microsoft to reduce the number and severity of vulnerabilities in our products and services. All Microsoft software development teams must follow SDL requirements, and we continuously update the SDL to reflect the changing threat landscape, industry best practices, and regulatory standards for compliance.

How does Microsoft's SDL improve application security?

The SDL process at Microsoft can be thought of in terms of five phases of development: requirements, design, implementation, verification, and release. It begins by defining software requirements with security in mind. To do meet this goal, we ask security-relevant questions about what the application must accomplish. Does the application need to collect sensitive data? Will the application perform sensitive or important tasks? Does the application need to accept input from untrusted sources?

Once relevant security requirements have been identified, we design our software to incorporate security features that meet these requirements. Our developers implement SDL and design requirements in the code, which we verify through manual code review, automated security tooling, and penetration testing. Finally, before code can be released, new features and material changes undergo final security and privacy review to ensure all requirements are met.

How does Microsoft test source code for common vulnerabilities?

To support our developers in implementing security requirements during code development and after release, Microsoft provides a suite of secure development tools to automatically check source code for security flaws and vulnerabilities. Microsoft defines and publishes a list of approved tools for our developers to use, such as compilers and development environments, along with the built-in security checks executed automatically within Microsoft build pipelines. Our developers use the latest versions of approved tools to take advantage of new security features.

Before code can be checked into a release branch, the SDL requires manual code review by a separate reviewer. Code reviewers check for coding errors and verify that code changes meet SDL and design requirements, pass functional and security tests, and perform reliably. They also review associated documentation, configs, and dependencies to ensure code changes are documented appropriately and will not cause unintended side-effects. If a reviewer finds problems during code review, they can ask the submitter to resubmit the code with suggested changes and additional testing. Code reviewers may also decide to block check-in entirely for code that does not meet requirements. Once the code has been deemed satisfactory by the reviewer, the reviewer provides approval, which is required before code can proceed to the next deployment phase.

In addition to secure development tools and manual code review, Microsoft enforces SDL requirements using automated security tooling. Many of these tools are built into the commit pipeline and automatically analyze code for security flaws as it is checked in and as new builds are compiled. Examples include static code analysis for common security flaws and credential scanners that analyze code for embedded secrets. Issues uncovered by automated security tools must be fixed before new builds can pass security review and be approved for release.

How does Microsoft manage open-source software?

Microsoft has adopted a high-level strategy for managing open-source security, which uses tools and workflows designed to:

- Understand which open-source components are being used in our products and services.
- Track where and how those components are used.
- Determine whether those components have any vulnerabilities.
- Respond properly when vulnerabilities are discovered that affect those components.

Microsoft engineering teams maintain responsibility for the security of all open-source software included in a product or service. To achieve this security at scale, Microsoft has built essential capabilities into engineering systems through Component Governance (CG), which automates open-source detection, legal requirement workflows, and alerting for vulnerable components. Automated CG tools scan builds at Microsoft for open-source components and associated security vulnerabilities or legal obligations. Discovered components are registered and submitted to the appropriate teams for business and security reviews. These reviews are designed to evaluate any legal obligations or security vulnerabilities associated with open-source components and resolve them before approving components for deployment.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to security development and operation.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
ISO 27001/27002 Statement of Applicability Certificate	A.12.1.2: Change management controls A.14.2: Security in development and support processes	April 24, 2023
ISO 27017 Statement of Applicability Certificate	A.12.1.2: Change management controls A.14.2: Security in development and support processes	April 24, 2023
SOC 1 SOC 2 SOC 3	SDL-1: Security Development Lifecycle (SDL) methodology SDL-2: Security control requirements documented in releases SDL-4: Segregation of test and production environments SDL-6: Malware scans on source code builds SDL7: Semi-annual SDL review	November 17, 2023

Microsoft 365

Expand table

External audits	Section	Latest report date
FedRAMP (Office 365)	SA-3: System Development Life Cycle	July 31, 2023
ISO 27001/27002/27017 Statement of Applicability Certification (27001/27002) Certification (27017)	A.12.1.2: Change management controls A.14.2: Security in development and support processes	March 2023
SOC 1 SOC 2	CA-03: Risk management CA-18: Change management CA-19: Change monitoring CA-21: Change testing CA-38: Baseline configurations CA-46: Security review	January 3, 2023

Resources

- Video: Microsoft 365 security development and operation
- Microsoft's Security Development Lifecycle

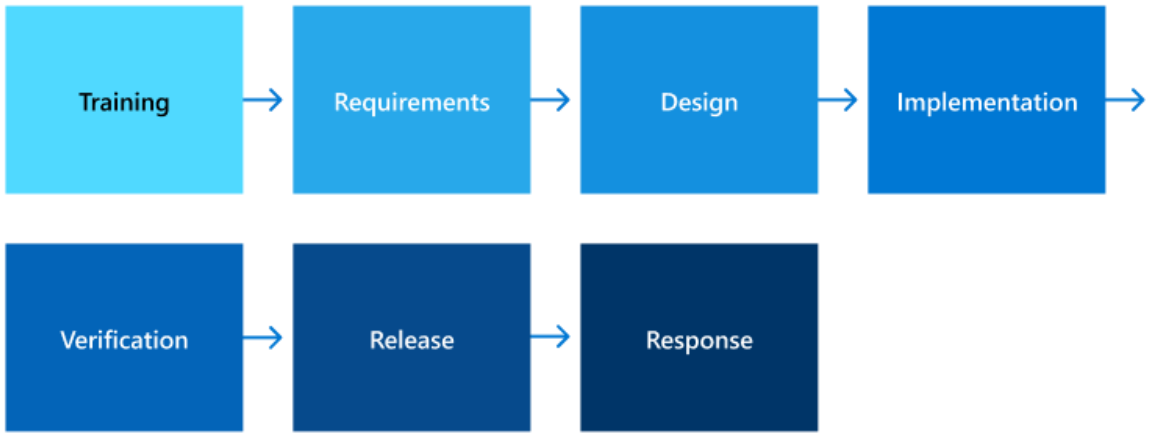
Feedback

Was this page helpful? Yes No

Microsoft Security Development Lifecycle (SDL)

Article • 03/02/2023

Security and privacy should never be an afterthought when developing secure software, a formal process must be in place to ensure they're considered at all points of the product's lifecycle. Microsoft's Security Development Lifecycle (SDL) embeds comprehensive security requirements, technology specific tooling, and mandatory processes into the development and operation of all software products. All development teams at Microsoft must adhere to the SDL processes and requirements, resulting in more secure software with fewer and less severe vulnerabilities at a reduced development cost.



Microsoft SDL consists of seven components including five core phases and two supporting security activities. The five core phases are requirements, design, implementation, verification, and release. Each of these phases contains mandatory checks and approvals to ensure all security and privacy requirements and best practices are properly addressed. The two supporting security activities, training and response, are conducted before and after the core phases respectively to ensure they're properly implemented, and software remains secure after deployment.

Training

All Microsoft employees are required to complete general security awareness training and specific training appropriate to their role. Initial security awareness training is provided to new employees upon hire and annual refresher training is required throughout their employment at Microsoft.

Developers and engineers must also participate in role specific training to keep them informed on security basics and recent trends in secure development. All full-time employees, interns, contingent staff, subcontractors, and third parties are also encouraged and provided with the opportunity to seek advanced security and privacy training.

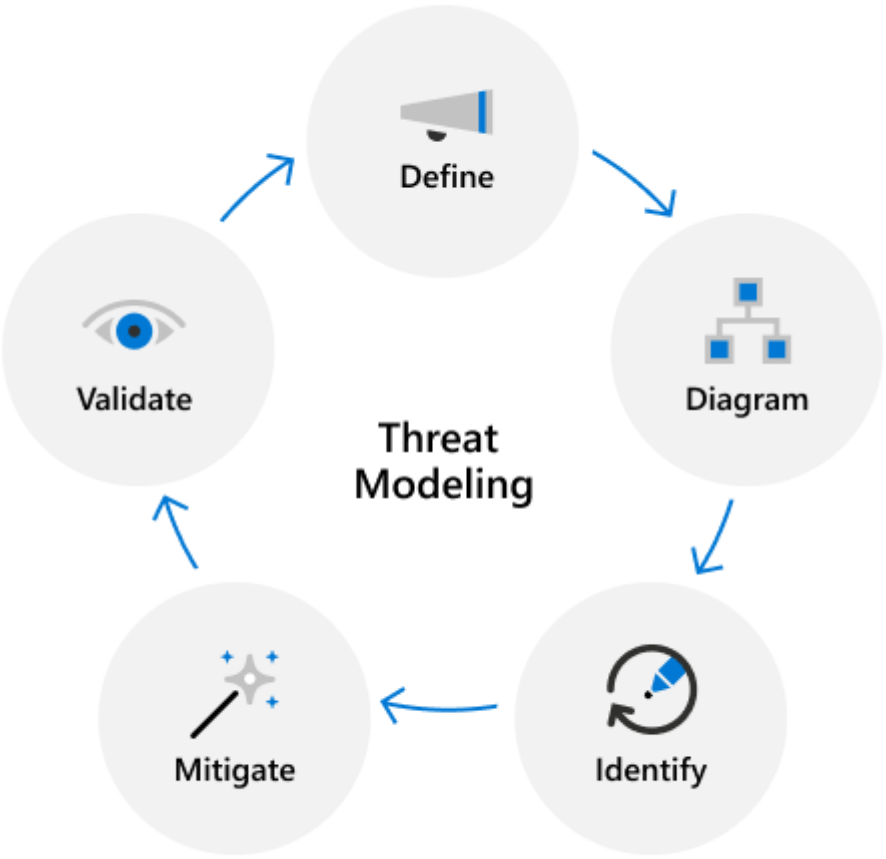
Requirements

Every product, service, and feature Microsoft develops starts with clearly defined security and privacy requirements; they're the foundation of secure applications and inform their design. Development teams define these requirements based on factors such as the type of data the product will handle, known threats, best practices, regulations and industry requirements, and lessons learned from previous incidents. Once defined, the requirements are clearly defined, documented, and tracked.

Software development is a continuous process, meaning that the associated security and privacy requirements change throughout the product's lifecycle to reflect changes in functionality and the threat landscape.

Design

Once the security, privacy, and functional requirements have been defined, the design of the software can begin. As a part of the design process, threat models are created to help identify, categorize, and rate potential threats according to risk. Threat models must be maintained and updated throughout the lifecycle of each product as changes are made to the software.



The threat modeling process begins by defining the different components of a product and how they interact with each other in key functional scenarios, such as authentication. Data Flow Diagrams (DFDs) are created to visually represent key data flow interactions, data types, ports, and protocols used. DFDs are used to identify and prioritize threats for mitigation that are added to the product's security requirements.

Developers are required to use Microsoft's Threat Modeling Tool for all threat models, which enables the team to:

- Communicate about the security design of their systems
- Analyze security designs for potential security issues using a proven methodology
- Suggest and manage mitigation for security issues

Before any product is released, all threat models are reviewed for accuracy and completeness, including mitigation for unacceptable risks.

Implementation

Implementation begins with developers writing code according to the plan they created in the previous two phases. Microsoft provides developers with a suite of secure development tools to effectively implement all the security, privacy, and function

requirements of the software they design. These tools include compilers, secure development environments, and built-in security checks.

Verification

Before any written code can be released, several checks and approvals are required to verify that the code conforms to SDL, meets design requirements, and is free of coding errors. SDL requires that manual reviews are conducted by a reviewer separate from the personnel the developed the code. Separation of duties is an important control in this step to ensure no code can be written and released by the same person leading to potential accidental or malicious harm.

Various automated checks are also required and are built into the commit pipeline to analyze code during check-in and when builds are compiled. The security checks used at Microsoft fall into the following categories:

- **Static code analysis**: Analyzes source code for potential security flaws, including the presence of credentials in code.
- **Binary analysis**: Assesses vulnerabilities at the binary code level to confirm code is production ready.
- **Credential and secret scanner**: Identify possible instances of credential and secret exposure in source code and configuration files.
- **Encryption scanning**: Validates encryption best practices in source code and code execution.
- **Fuzz testing**: Use malformed and unexpected data to exercise APIs and parsers to check for vulnerabilities and validate error handling.
- **Configuration validation**: Analyzes the configuration of production systems against security standards and best practices.
- **Component Governance (CG)**: Open-source software detection and checking of version, vulnerability, and legal obligations.

If the manual reviewer or automated tools find any issues in the code, the submitter will be notified, and they're required to make the necessary changes before submitting it for review again.

Additionally, penetration tests are regularly conducted on Microsoft online services by both internal and external providers. Penetration tests provide another means for discovering security flaws not detected by other methods. To learn more about penetration testing at Microsoft, see Attack simulation in Microsoft 365.

Release

After passing all required security tests and reviews, builds aren't immediately released to all customers. Builds are systematically and gradually released to larger and larger groups, referred to as rings, in what is called a safe deployment process (SDP). The SDP rings are defined as follows:

- **Ring 0**: The development team responsible for the service
- **Ring 1**: All Microsoft employees
- **Ring 2**: Users outside of Microsoft who have configured their organization or specific users to be on the targeted release channel
- **Ring 3**: Worldwide standard release in sub-phases

Builds remain in each of these rings for an appropriate number of days with high load periods, except for Ring 3 since the build has been appropriately tested for stability in the earlier rings.

Response

All Microsoft services are extensively logged and monitored after release, identifying potential security incidents using a centralized proprietary near-real-time monitoring system. To learn more about security monitoring and security incident management at Microsoft, see Security monitoring overview and Microsoft security incident management.

What is Responsible AI?

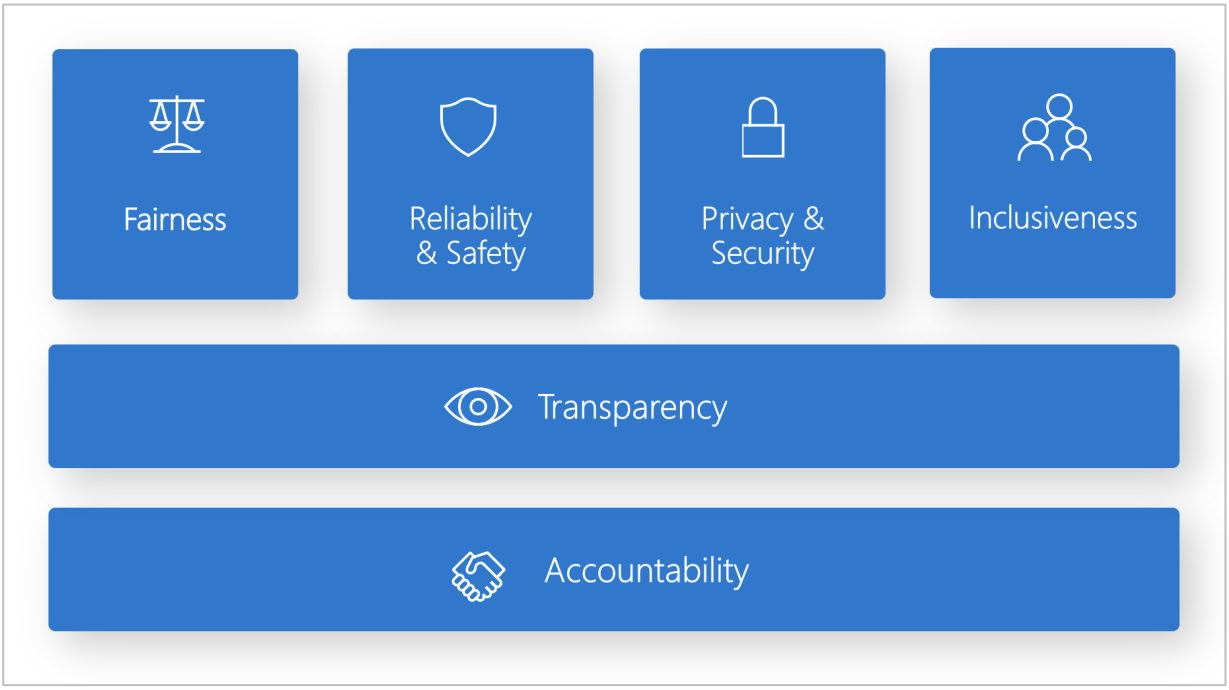
Article • 11/09/2022

APPLIES TO: Azure CLI ml extension v2 (current) Python SDK azure-ai-ml v2 (current)

Responsible Artificial Intelligence (Responsible AI) is an approach to developing, assessing, and deploying AI systems in a safe, trustworthy, and ethical way. AI systems are the product of many decisions made by those who develop and deploy them. From system purpose to how people interact with AI systems, Responsible AI can help proactively guide these decisions toward more beneficial and equitable outcomes. That means keeping people and their goals at the center of system design decisions and respecting enduring values like fairness, reliability, and transparency.

Microsoft has developed a Responsible AI Standard. It's a framework for building AI systems according to six principles: fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability. For Microsoft, these principles are the cornerstone of a responsible and trustworthy approach to AI, especially as intelligent technology becomes more prevalent in products and services that people use every day.

This article demonstrates how Azure Machine Learning supports tools for enabling developers and data scientists to implement and operationalize the six principles.



Fairness and inclusiveness

AI systems should treat everyone fairly and avoid affecting similarly situated groups of people in different ways. For example, when AI systems provide guidance on medical treatment, loan applications, or employment, they should make the same recommendations to everyone who has similar symptoms, financial circumstances, or professional qualifications.

Fairness and inclusiveness in Azure Machine Learning: The fairness assessment component of the Responsible AI dashboard enables data scientists and developers to assess model fairness across sensitive groups defined in terms of gender, ethnicity, age, and other characteristics.

Reliability and safety

To build trust, it's critical that AI systems operate reliably, safely, and consistently. These systems should be able to operate as they were originally designed, respond safely to unanticipated conditions, and resist harmful manipulation. How they behave and the variety of conditions they can handle reflect the range of situations and circumstances that developers anticipated during design and testing.

Reliability and safety in Azure Machine Learning: The error analysis component of the Responsible AI dashboard enables data scientists and developers to:

- Get a deep understanding of how failure is distributed for a model.
- Identify cohorts (subsets) of data with a higher error rate than the overall benchmark.

These discrepancies might occur when the system or model underperforms for specific demographic groups or for infrequently observed input conditions in the training data.

Transparency

When AI systems help inform decisions that have tremendous impacts on people's lives, it's critical that people understand how those decisions were made. For example, a bank might use an AI system to decide whether a person is creditworthy. A company might use an AI system to determine the most qualified candidates to hire.

A crucial part of transparency is *interpretability*: the useful explanation of the behavior of AI systems and their components. Improving interpretability requires stakeholders to comprehend how and why AI systems function the way they do. The stakeholders can then identify potential performance issues, fairness issues, exclusionary practices, or unintended outcomes.

Transparency in Azure Machine Learning: The model interpretability and counterfactual what-if components of the Responsible AI dashboard enable data scientists and developers to generate human-understandable descriptions of the predictions of a model.

The model interpretability component provides multiple views into a model's behavior:

- *Global explanations*. For example, what features affect the overall behavior of a loan allocation model?
- *Local explanations*. For example, why was a customer's loan application approved or rejected?
- *Model explanations for a selected cohort of data points*. For example, what features affect the overall behavior of a loan allocation model for low-income applicants?

The counterfactual what-if component enables understanding and debugging a machine learning model in terms of how it reacts to feature changes and perturbations.

Azure Machine Learning also supports a Responsible AI scorecard. The scorecard is a customizable PDF report that developers can easily configure, generate, download, and share with their technical and non-technical stakeholders to educate them about their datasets and models health, achieve compliance, and build trust. This scorecard can also be used in audit reviews to uncover the characteristics of machine learning models.

Privacy and security

As AI becomes more prevalent, protecting privacy and securing personal and business information are becoming more important and complex. With AI, privacy and data security require close attention because access to data is essential for AI systems to make accurate and informed predictions and decisions about people. AI systems must comply with privacy laws that:

- Require transparency about the collection, use, and storage of data.
- Mandate that consumers have appropriate controls to choose how their data is used.

Privacy and security in Azure Machine Learning: Azure Machine Learning enables administrators and developers to create a secure configuration that complies with their companies' policies. With Azure Machine Learning and the Azure platform, users can:

- Restrict access to resources and operations by user account or group.
- Restrict incoming and outgoing network communications.
- Encrypt data in transit and at rest.
- Scan for vulnerabilities.

- Apply and audit configuration policies.

Microsoft has also created two open-source packages that can enable further implementation of privacy and security principles:

- SmartNoise: Differential privacy is a set of systems and practices that help keep the data of individuals safe and private. In machine learning solutions, differential privacy might be required for regulatory compliance. SmartNoise is an open-source project (co-developed by Microsoft) that contains components for building differentially private systems that are global.

- Counterfit: Counterfit is an open-source project that comprises a command-line tool and generic automation layer to allow developers to simulate cyberattacks against AI systems. Anyone can download the tool and deploy it through Azure Cloud Shell to run in a browser, or deploy it locally in an Anaconda Python environment. It can assess AI models hosted in various cloud environments, on-premises, or in the edge. The tool is agnostic to AI models and supports various data types, including text, images, or generic input.

Accountability

The people who design and deploy AI systems must be accountable for how their systems operate. Organizations should draw upon industry standards to develop accountability norms. These norms can ensure that AI systems aren't the final authority on any decision that affects people's lives. They can also ensure that humans maintain meaningful control over otherwise highly autonomous AI systems.

Accountability in Azure Machine Learning: Machine learning operations (MLOps) is based on DevOps principles and practices that increase the efficiency of AI workflows. Azure Machine Learning provides the following MLOps capabilities for better accountability of your AI systems:

- Register, package, and deploy models from anywhere. You can also track the associated metadata that's required to use the model.
- Capture the governance data for the end-to-end machine learning lifecycle. The logged lineage information can include who is publishing models, why changes were made, and when models were deployed or used in production.
- Notify and alert on events in the machine learning lifecycle. Examples include experiment completion, model registration, model deployment, and data drift detection.
- Monitor applications for operational issues and issues related to machine learning. Compare model inputs between training and inference, explore model-specific

metrics, and provide monitoring and alerts on your machine learning infrastructure.

Besides the MLOps capabilities, the Responsible AI scorecard in Azure Machine Learning creates accountability by enabling cross-stakeholder communications. The scorecard also creates accountability by empowering developers to configure, download, and share their model health insights with their technical and non-technical stakeholders about AI data and model health. Sharing these insights can help build trust.

The machine learning platform also enables decision-making by informing business decisions through:

- Data-driven insights, to help stakeholders understand causal treatment effects on an outcome, by using historical data only. For example, "How would a medicine affect a patient's blood pressure?" These insights are provided through the causal inference component of the Responsible AI dashboard.
- Model-driven insights, to answer users' questions (such as "What can I do to get a different outcome from your AI next time?") so they can take action. Such insights are provided to data scientists through the counterfactual what-if component of the Responsible AI dashboard.

Next steps

- For more information on how to implement Responsible AI in Azure Machine Learning, see Responsible AI dashboard.
- Learn how to generate the Responsible AI dashboard via CLI and SDK or Azure Machine Learning studio UI.
- Learn how to generate a Responsible AI scorecard based on the insights observed in your Responsible AI dashboard.
- Learn about the Responsible AI Standard for building AI systems according to six key principles.

Platform code integrity

Article • 11/11/2022

A significant challenge in operating a complex system like Microsoft Azure is ensuring that only authorized software is running in the system. Unauthorized software presents several risks to any business:

- Security risks such as dedicated attack tools, custom malware, and third-party software with known vulnerabilities
- Compliance risks when the approved change management process isn't used to bring in new software
- Quality risk from externally developed software, which may not meet the operational requirements of the business

In Azure, we face the same challenge and at significant complexity. We have thousands of servers running software developed and maintained by thousands of engineers. This presents a large attack surface that cannot be managed through business processes alone.

Adding an authorization gate

Azure uses a rich engineering process that implements gates on the security, compliance, and quality of the software we deploy. This process includes access control to source code, conducting peer code reviews, doing static analysis for security vulnerabilities, following Microsoft's Security Development Lifecycle (SDL), and conducting functional and quality testing. We need to guarantee that the software we deploy has flowed through this process. Code integrity helps us achieve that guarantee.

Code integrity as an authorization gate

Code integrity is a kernel level service that became available starting in Windows Server 2016. Code integrity can apply a strict execution control policy whenever a driver or a dynamically linked library (DLL) is loaded, an executable binary is executed, or a script is run. Similar systems, such as DM-Verity, exist for Linux. A code integrity policy consists of a set of authorization indicators, either code signing certificates or SHA256 file hashes, which the kernel matches before loading or executing a binary or script.

Code Integrity allows a system administrator to define a policy that authorizes only binaries and scripts that have been signed by particular certificates or match specified

SHA256 hashes. The kernel enforces this policy by blocking execution of everything that doesn't meet the set policy.

A concern with a code integrity policy is that unless the policy is perfectly correct, it can block critical software in production and cause an outage. Given this concern, one may ask why it isn't sufficient to use security monitoring to detect when unauthorized software has executed. Code integrity has an audit mode that, instead of preventing execution, can alert when unauthorized software is run. Alerting certainly can add much value in addressing compliance risks, but for security risks such as ransomware or custom malware, delaying the response by even a few seconds can be the difference between protection and an adversary gaining a persistent foothold in your fleet. In Azure, we've invested significantly to manage any risk of code integrity contributing to a customer impacting outage.

Build process

As discussed above, the Azure build system has a rich set of tests to ensure software changes are secure and compliant. Once a build has progressed through validation, the build system signs it using an Azure build certificate. The certificate indicates the build has passed through the entire change management process. The final test that the build goes through is called Code Signature Validation (CSV). CSV confirms the newly built binaries meet the code integrity policy before we deploy to production. This gives us high confidence that we won't cause a customer impacting outage because of incorrectly signed binaries. If CSV finds a problem, the build breaks and the relevant engineers are paged to investigate and fix the issue.

Safety during deployment

Even though we perform CSV for every build, there's still a chance that some change or inconsistency in production may cause a code integrity related outage. For example, a machine may be running an old version of the code integrity policy or it may be in an unhealthy state that produces false positives in code integrity. (At Azure scale, we've seen it all.) As such, we need to continue to protect against the risk of an outage during deployment.

All changes in Azure are required to deploy through a series of stages. The first of these are internal Azure testing instances. The next stage is used only to serve other Microsoft product teams. The final stage serves third-party customers. When a change is deployed, it moves to each of these stages in turn, and pauses to measure the health of the stage. If the change is found to have no negative impact, then it moves to the next

stage. If we make a bad change to a code integrity policy, the change is detected during this staged deployment and rolled back.

Incident response

Even with this layered protection, it's still possible that some server in the fleet may block properly authorized software and cause a customer facing issue, one of our worst-case scenarios. Our final layer of defense is human investigation. Each time code integrity blocks a file, it raises an alert for the on-call engineers to investigate. The alert allows us to start security investigations and intervene, whether the issue is an indicator of a real attack, a false positive, or other customer-impacting situation. This minimizes the time it takes to mitigate any code integrity related issues.

Next steps

Learn how Windows 10 uses configurable code integrity.

To learn more about what we do to drive platform integrity and security, see:

- Firmware security
- Secure boot
- Measured boot and host attestation
- Project Cerberus
- Encryption at rest
- Hypervisor security

Microsoft 365 change management

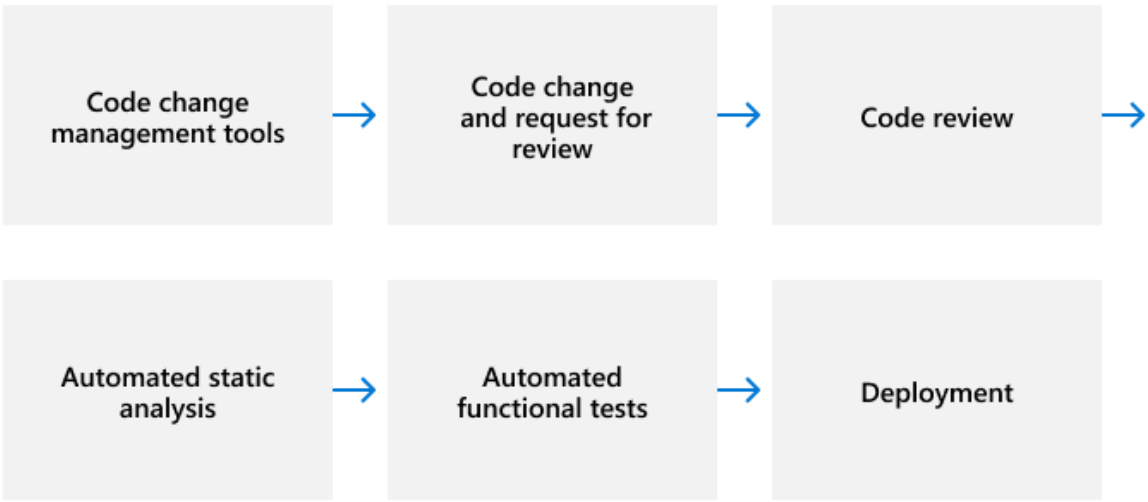
Article • 03/02/2023

Microsoft 365 enforces change management procedures when both code and non-code changes to its systems are made to maintain its security posture. Any configuration drift from the initial posture can introduce vulnerabilities, break functionality, or disrupt availability. Once an information system within Microsoft 365 has been deployed with a robust security posture, detailed change management processes are enforced to maintain system integrity.

There are many drivers of change in Microsoft 365, including new functional or security requirements, feedback from customers, identified vulnerabilities, and audit findings. Regardless of the driver for change, service teams use ticketing or source control tools to document evidence of approval and track all changes.

Source code changes

Changes are deployed through Microsoft's Secure Development Lifecycle (SDL), which is followed by all engineering and development projects in Microsoft 365. This is a software development model that includes specific security considerations related to code reviews, tests, and approvals before they're systematically released into the Microsoft 365 environment.



The SDL acts as a framework and includes the identification of possible risks to the finished development project and mitigation strategies that can be implemented and tested during the development phases. Critical security review and approval checkpoints are included as well.

Change identification and planning

Service teams meet regularly to discuss proposed changes, including justification, scope, security impact, priority, dependencies, deployment plans, roles, and responsibilities. This information is documented in the change management tracking system. If the change is rejected, the justification is explicitly documented in the ticket for future reference.

Personnel code reviews

Developers tasked with implementing a code change submit a pull request that replicates the main branch's code, allowing them to make necessary modifications. Before any new code can be included in a new build and deployed, it must pass personnel code review. Enforcement of these reviews is handled through an automated code pipeline attached to each code repository and can't be circumvented. Once the required approval is received, the code can move on to the next phase.

Code reviewers check for coding errors, verify that the changes meet the requirements, and perform a security impact analysis. Reviews must be conducted by someone other than the people who developed the code, enforcing the principle of separation of duties. Preventing the same people from submitting and approving their own code is a critical control that Microsoft strictly enforces. This greatly reduces the possibility of people single-handedly releasing, either intentionally or unintentionally, harmful, or buggy code. If reviewers find problems during the code review, they halt the change, and have developers resubmit the code with suggested changes and additional testing. Code reviewers may also decide to reject check-in entirely for code that doesn't meet the identified requirements. Once the code is deemed satisfactory by the reviewer, approval is provided, and the code is checked into the main branch as a commit.

Automated build pipeline and security checks

Once all changes for the sprint are committed to the main branch, the automated build process begins. This is where the code is subjected to various automated security checks. These checks include static code analysis, binary analysis, and encryption scanning. Microsoft 365 defines a set of essential tests that each build must pass prior to deployment to pre-production environments. Builds that don't pass are rejected and sent back to the development team where the necessary adjustments are made until they can reach the threshold pass rate. Successful builds proceed to the pre-production environment via an automated deployment pipeline.

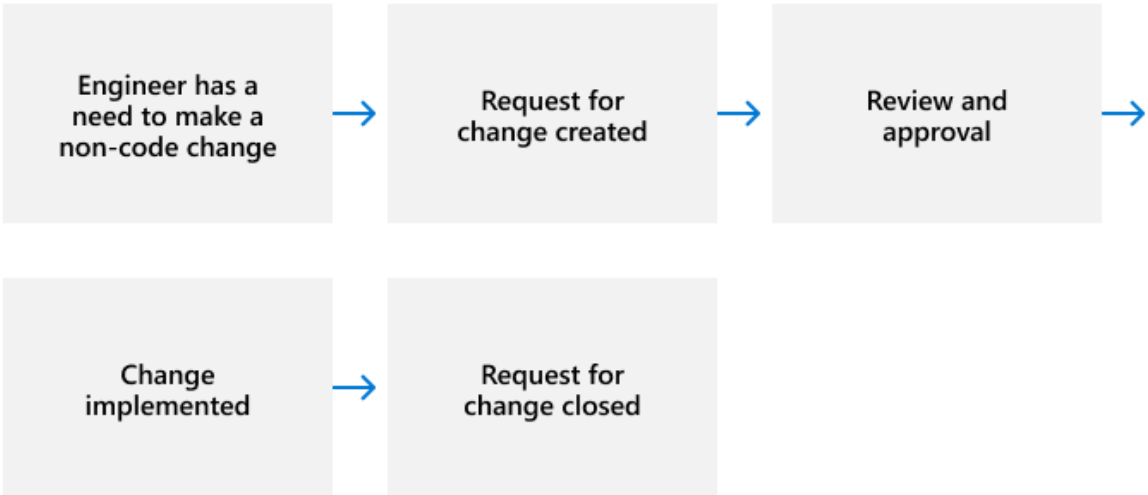
Build release

Builds are initially released to only the service team that developed the feature. It must function without issue before being released to progressively larger test groups in logically isolated cloud environments called rings. After the service team, the build is released to all internal Microsoft 365 groups, followed by release of the build to all internal Microsoft groups. This testing, often referred to internally as *dogfooding*, allows Microsoft to identify bugs in the true production environment prior to the build being released to external customers. These testing methods ensure Microsoft's code is secure and functioning as expected before it reaches customers and worldwide deployment. Previous builds are always retained for rollback purposes.

The engineering teams determine the amount of time a build spends in each ring during high load periods before proceeding to the next ring. If all testing is successful in each internal ring, the build is released to customers worldwide, first as a Targeted Release to customer tenants who have opted into that ring, followed by a Worldwide Standard Release.

Non-code changes

Non-code changes are defined as any modifications to Microsoft 365 systems that don't involve creating or editing service source code. This can include the opening of ports, changing of Access Control Lists (ACLs), or other changes to the underlying system. In comparison, non-code changes occur less frequently than code changes but still require a high level of scrutiny.



A description of the change is documented along with implementation steps, validation steps, and a rollback plan. Before the change is implemented, the plans are peer

reviewed for accuracy and security impact by at least one person. Once approved, the documented plans are implemented. If all validation steps successfully pass, the results are documented in the ticket, and it's marked as resolved.

If the implementation of the change is unsuccessful, the rollback plans are triggered, and the team returns to the planning phase and repeats the process until successful.

Resources

- Microsoft 365 change guide - Deploy Office

Microsoft 365 service compliance program

Article • 04/19/2023

Microsoft runs on trust and Microsoft 365 customers require world class security and privacy functions within their products. The thorough auditing of Microsoft 365 assures customers that the security functions, data handling practices, associated policies, and tooling work in harmony to meet customer needs for certifications, accreditations, and attestations. Microsoft 365's internal compliance program is designed to ensure security and privacy are considered at all phases of the development process.

Each service begins the assurance lifecycle with the execution of three related efforts: security, privacy, and compliance. Microsoft directs these efforts by using assessments. The security assessment ensures that all security tools, application policies, and functions are applied and appropriately configured. The privacy assessment reviews all data a service collects, transmits, and stores. It confirms strict data handling is in place and that it adheres to regional privacy requirements. Finally, the compliance assessment verifies that the services satisfy all the relevant compliance requirements for the customers they wish to serve.

Let's take a deeper dive into each of these components to understand their role in the overall process.

Security

The security assessment is owned by the Microsoft 365 Security Team and is designed to identify all the security solutions that service teams need to install and configure. Microsoft 365 has developed solutions for service teams to secure each individual service and connect them to the centralized systems that helps protect Microsoft as a whole. These solutions include identity and access management, anti-malware software, central logging, TLS configuration, and others. Service teams are also required to create data flow diagrams for threat modeling to help map potential attack vectors. The Microsoft 365 Security Team is available to service teams for guidance and perform a final review and approval of the service security functions.

Privacy

Privacy focuses on customer data the service team transmits, processes, and stores. Tasks include identifying data types collected, retention periods, classification, and any

third-party interactions. Once service teams complete their tasks, a dedicated privacy manager and a lawyer perform a review of the data handling functions. Threat models are also used to safeguard against unexpected data leakage. The privacy manager must provide official approval and verification of privacy assessment completion.

Compliance

Many customers have industry specific compliance requirements met to use cloud services such as Microsoft 365. An objective of compliance assessments is to ensure that the service and any downstream dependencies meet the compliance requirements for applicable accreditations. These accreditations may include the ISO 27001 Information security standard, SOC II Type 2, FedRAMP for government customers, and others.

The final step of the compliance assessment is the Final Trust Review conducted by Microsoft's Trust Architects. This provides an all-encompassing check of a service's security, privacy, and compliance posture. It verifies each service is properly secured, follows best practices, meets all relevant regulatory requirements, and all identifiable risks have been adequately addressed. Additionally, the holistic relationship to all other services and environments is considered at this point.

Teams developing a new service are encouraged to participate in a consultative review session with the Trust Architects in combination with other assessments. The service teams receive advice on their design, including guidance on avoiding potential pitfalls and blockers prior to beginning development.

Once the security, privacy, compliance efforts are complete and approved, a service may become generally available to customers or need to wait for audit depending on the size, scope, and service.

Security incident management overview

Article • 08/08/2023

What is a security incident?

Microsoft defines a security incident in its online services as a confirmed breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorized disclosure of, or access to customer data or personal data while being processed by Microsoft. For example, unauthorized access to Microsoft online services infrastructure and exfiltration of customer data would constitute a security incident, while compliance events that don't affect the confidentiality, integrity, or availability of services or customer data aren't considered security incidents.

How does Microsoft respond to security incidents?

Whenever there's a security incident, Microsoft strives to respond quickly and effectively to protect Microsoft services and customer data. Microsoft employs an incident response strategy designed to investigate, contain, and remove security threats quickly and efficiently.

Microsoft cloud services are continuously monitored for signs of compromise. In addition to automated security monitoring and alerting, all employees receive annual training to recognize and report signs of potential security incidents. Any suspicious activity detected by employees, customers, or security monitoring tools are escalated to Service-specific Security Response teams for investigation. All service operations teams, including Service-specific Security Response teams, maintain a deep on-call rotation to ensure resources are available for incident response 24x7x365. Our on-call rotations enable Microsoft to mount an effective incident response at any time or scale, including widespread or concurrent events.

When suspicious activity is detected and escalated, Service-specific Security Response teams initiate a process of **analysis, containment, eradication, and recovery**. These teams coordinate analysis of the potential incident to determine its scope, including any impact to customers or customer data. Based on this analysis, Service-specific Security Response teams work with impacted service teams to develop a plan to contain the threat and minimize the impact of the incident, eradicate the threat from the environment, and fully recover to a known secure state. Relevant service teams implement the plan with support from Service-specific Security Response teams to

ensure the threat is successfully eliminated and impacted services undergo a complete recovery.

After an incident is resolved, service teams implement any lessons learned from the incident to better prevent, detect, and respond to similar incidents in the future. Select security incidents, especially those incidents that are customer-impacting or result in a data breach, undergo a full incident post-mortem. The post-mortem is designed to identify technical lapses, procedural failures, manual errors, and other process flaws that might have contributed to the incident or that were identified during the incident response process. Improvements identified during the post-mortem are implemented with coordination from Service-specific Security Response teams to help prevent future incidents and improve detection and response capabilities.

How and when are customers notified of security or privacy incidents?

Whenever Microsoft becomes aware of a breach of security involving unauthorized loss, disclosure, or modification of customer data, Microsoft notifies affected customers within 72 hours as outlined in the Data Protection Addendum (DPA). The notification timeline commitment begins when the official security incident declaration occurs. Upon declaring a security incident, the notification process occurs as expeditiously as possible, without undue delay.

Notifications include a description of the nature of the breach, approximate user impact, and mitigation steps (if applicable). If Microsoft's investigation isn't complete at the time of initial notification, the notification will also indicate next steps and timelines for subsequent communication.

If a customer becomes aware of an incident that could have an impact on Microsoft, including but not limited to a data breach, the customer is responsible for promptly notifying Microsoft of the incident as defined in the DPA.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to incident management.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
ISO 27001/27002 Statement of Applicability Certificate	A.16.1: Management of information security incidents and improvements	April 24, 2023
ISO 27017 Statement of Applicability Certificate	A.16.1: Management of information security incidents and improvements	April 24, 2023
ISO 27018 Statement of Applicability Certificate	A.9.1: Notification of a data breach involving PII	April 24, 2023
SOC 1	IM-1: Incident management framework IM-2: Detection mechanisms and alerting IM-3: Incident response execution IM-4: Incident postmortems IM-6: Incident response testing OA-7: On-call engineer access	November 17, 2023
SOC 2 SOC 3	CCM-9: Forensic procedures CUEC: Reporting incidents IM-1: Incident management framework IM-2: Detection mechanisms and alerting IM-3: Incident response execution IM-4: Incident postmortems IM-6: Incident response testing OA-7: On-call engineer access SOC2-6: Customer Support website SOC2-9: Service dashboards	November 17, 2023

Microsoft 365

Expand table

External audits	Section	Latest report date
FedRAMP (Office 365)	IR-4: Incident handling IR-6: Incident reporting IR-8: Incident response plan	July 31, 2023
ISO 27001/27002/27017 Statement of Applicability Certification (27001/27002) Certification (27017)	A.16.1: Management of information security incidents and improvements	March 2023
ISO 27018 Statement of Applicability Certificate	A.10.1: Notification of a data breach involving PII	March 2023
SOC 1	CA-26: Security incident reporting CA-47: Incident response	January 3, 2023
SOC 2	CA-12: Service level agreements (SLAs) CA-13: Incident response guides CA-15: Service health notifications CA-26: Security incident reporting CA-29: On-call engineers CA-47: Incident response	January 3, 2023
SOC 3	CUEC-08: Reporting incidents	January 3, 2023

Resources

- Video: Microsoft Online Services incident management
- Product Terms
- Data Protection Addendum (DPA)
- Office 365 - Third-Party Vulnerability Assessment - 2022

Feedback

Was this page helpful?

Yes No

Microsoft security incident management

Article • 03/02/2023

Microsoft works continuously to provide highly secure, enterprise-grade services for Microsoft customers, but security incidents are an inevitable reality that must be thoroughly and swiftly managed. This document provides an overview on how Microsoft handles security incidents using tried and true methods and technologies to minimize their potential impact. A security incident refers to any unlawful access to customer data stored on Microsoft's equipment or in Microsoft's facilities, or unauthorized access to such equipment or facilities that have the potential to result in the loss, disclosure, or alteration of customer data. Microsoft's goals when responding to security incidents are to protect customer data and Microsoft's online services.

Microsoft online services security teams and the various service teams work jointly and take the same approach to security incidents:

- Preparation
- Detection and Analysis
- Containment, Eradication, and Recovery
- Post-Incident Activity

Microsoft approach to security incident management

Microsoft's approach to managing a security incident conforms to National Institute of Standards and Technology (NIST) Special Publication (SP) 800-61. Microsoft has several dedicated teams that work together to prevent, monitor, detect, and respond to security incidents.

Team/Area	Description
Microsoft Security Response Center	Identifies, monitors, resolves, and responds to security incidents and Microsoft software security vulnerabilities.
Cyber Defense Operations Center	The Cyber Defense Operations Center is the physical location that brings together security response teams and experts from across the company to help protect, detect, and respond to threats in real time.

Team/Area	Description
Corporate, External, and Legal Affairs	Provides legal and regulatory advice for a suspected security incident.
Microsoft Datacenter Security Team	Team that focuses across the various services on common security engineering investments to protect, detect, and respond to service architecture risks and threats.
Microsoft security response teams	Independent Azure, Dynamics 365, and Microsoft 365 security teams that partner with service teams to build the appropriate security incident management process and to drive any security incident response.
Microsoft Governance, Risk, and Compliance (GRC) teams	Provide guidance on regulatory requirements, compliance, and privacy.
Service teams	Engineering teams for Azure, Dynamics 365, Microsoft 365 that are responsible for security-related policies and decisions for each service.
Azure operations managers	Oversees the investigation and resolution of Azure-related security and privacy incidents.
Microsoft Threat Intelligence Center (MSTIC)	Provides the current state of art in digital security threats against Microsoft infrastructure and assets, helps partner teams inside Microsoft prioritize mitigation and prevention effort action plans, and increases protection by adopting near real-time incident monitoring/detection.
Customer experience communication teams	Engineering teams responsible for all customer communications about security and service incidents. Separate teams are dedicated to Azure, Dynamics 365, and Microsoft 365.

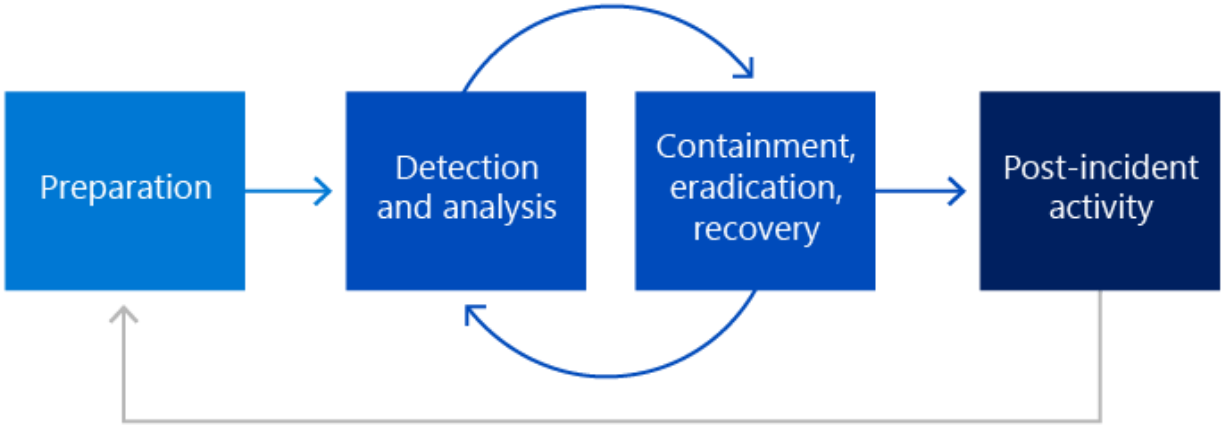
Response management process

Microsoft online services security teams and service teams work together on and take the same approach to security incidents, which is based on the NIST 800-61 response management phases:

- **Preparation**: Refers to the organizational preparation that is needed to be able to respond, including tools, processes, competencies, and readiness.
- **Detection & analysis**: Refers to the activity to detect a security incident in a production environment and to analyze all events to confirm the authenticity of

the security incident.

- **Containment, eradication, recovery**: Refers to the required and appropriate actions taken to contain the security incident based on the analysis done in the previous phase. More analysis may also be necessary in this phase to fully recovery from the security incident.
- **Post-incident activity**: Refers to the post-mortem analysis performed after the recovery of a security incident. The operational actions performed during the process are reviewed to determine if any changes need to be made in the preparation or detection and analysis phases.



Federated security response model

Microsoft online services consist of core Microsoft products, including Azure, Dynamics 365, and Microsoft 365. Each of these services are operated by separate teams with their own security operational processes. Other teams at Microsoft, such as MSTIC, are also engaged in various security aspects of Microsoft online services. Because of the multitude of teams working on security operations management across all the various services that make up Microsoft online services, Microsoft has implemented a federated security response model.

This table presents the operational boundaries between the various Microsoft online service security operations teams and the Microsoft service teams:

Activity	Microsoft Security Team Operations	Microsoft Service Team Operations

Activity	Microsoft Security Team Operations	Microsoft Service Team Operations
Detection and analysis	- Detection requirements - Security monitoring and analysis - Indicator of compromise (IOC) sweeps - Breach hunt - 24x7 security on-call and incident response lead	- Detection requirements - Monitoring infrastructure deployment - Service analysis and insight - Event and alert triage - 24x7 service engineering on-call
Containment, eradication, recovery	- Incident response lead - Forensics investigation - Security expertise and consulting - Recovery guidance	- Security incident owner - Service insight and expertise - Execute containment, eradication, and recovery
Post-incident activity	- Post-incident analysis lead - Data collection and archival - Lessons learned and bug requests - Incident reporting	- Service-side incident analysis - Prioritize follow-up activities - Implementation security investments - Service security readiness

Related articles

- Microsoft security incident management: Preparation
- Microsoft security incident management: Detection and analysis
- Microsoft security incident management: Containment, eradication, and recovery
- Microsoft security incident management: Post-incident activity

Microsoft security incident management: Preparation

Article • 04/04/2023

Training and background checks

Each employee working on Microsoft online services is provided with training regarding security incidents and response procedures that are appropriate to their role. Every Microsoft employee receives training upon joining, and annual refresher training every year thereafter. The training is designed to provide the employee with a basic understanding of Microsoft's approach to security so that upon completion of training, all employees understand:

- The definition of a security incident
- The responsibility of all employees to report security incidents
- How to escalate a potential security incident to the appropriate Microsoft security response team
- How Microsoft security incident response teams respond to security incidents
- Special concerns regarding privacy, particularly customer privacy
- Where to find additional information about security and privacy, and escalation contacts
- Any other relevant security areas (as needed)

The appropriate employees receive refresher training on security annually. The annual refresher training focuses on:

- Any changes made to the standard operating procedures in the preceding year
- The responsibility of everyone to report security incidents, and how to do so
- Where to find additional information about security and privacy, and escalation contacts
- Any other security focus areas that may be relevant each year

Each employee working on Microsoft online services is subject to an appropriate and thorough background check that includes the candidate's education, employment, criminal history, and other specific information per United States regulations like Health Insurance Portability and Accountability Act (HIPAA), International Traffic in Arms Regulations (ITAR), Federal Risk and Authorization Management Program (FedRAMP), and others.

The background checks are mandatory for all employees working within Microsoft engineering. Some Microsoft online service environments and operator roles may also require full fingerprinting, citizenship requirements, government clearance requirements, and other more stringent controls. In addition, some service teams and roles may go through specialized security training as needed. Lastly, the security team members themselves get specialized training and conference participation that relates directly to security. This training varies by need of the team and employee, but includes things like industry conferences, internal Microsoft Security conferences, and external training courses through well-known security training vendors in the industry. We also have dedicated security training articles published throughout the year for the security community across Microsoft and specialized to Microsoft online services regularly.

Penetration testing & assessment

Microsoft works with various industry bodies and security experts to understand new threats and evolving trends. Microsoft continuously assesses its own systems for vulnerabilities, and contracts with various independent, external experts who do the same.

The tests carried out for service hardening within Microsoft online services can be grouped into four general categories:

1. **Automated security testing:** Internal and external personnel regularly scan Microsoft online service environments based on Microsoft SDL practices, Open Web Application Security Project (OWASP) Top 10 risks, and emerging threats reported by different industry bodies.
2. **Vulnerability assessments:** Formal engagements with independent, third-party testers regularly validate whether key logical controls are operating effectively to fulfill the service obligations of various regulatory bodies. The assessments are carried out by Council of Registered Ethical Security Testers (CREST)-certified personnel and are based on OWASP Top 10 risks and other service-applicable threats. All threats found are tracked to closure.
3. **Continuous system vulnerability testing:** Microsoft carries out regular testing in which teams attempt to breach the system using emerging threats, blended threats, and/or advanced persistent threats, while other teams attempt to block such attempts to breach.
4. **Microsoft Online Services Bug Bounty Program**: This program operates a policy of allowing limited, customer-originated, vulnerability assessments on Microsoft online services. For more information, see Microsoft Online Services Bug Bounty Terms .

The Microsoft online services engineering teams periodically publish various compliance documents. Several of those documents are available under a non-disclosure agreement from the Microsoft Cloud Service Trust Portal or from the Service Assurance area of the Microsoft Purview compliance portal

Attack Simulation

Microsoft engages in ongoing attack simulation exercises and live-site penetration testing of our security and response plans with the intent to improve detection and response capability. Microsoft regularly simulates real-world breaches, conducts continuous security monitoring, and practices security incident response to validate and improve the security of Microsoft online services.

Microsoft executes an assume breach security strategy using two core teams:

Red teams

Microsoft Red teams are groups of full-time staff within Microsoft that focus on breaching Microsoft's infrastructure, platform, and Microsoft's own tenants and applications. They are the dedicated adversary (a group of ethical hackers) performing targeted and persistent attacks against online services (but not customer applications or data). They provide continuous 'full spectrum' validation (for example, technical controls, paper policy, human response, etc.) of service incident response capabilities.

Blue teams

Microsoft Blue teams are composed of dedicated sets of security responders, and members from across the security incident response, engineering, and operations teams. They are independent and operate separately from the Red teams. The Blue teams follow established security processes and use the latest tools and technologies to detect and respond to attacks and penetration attempts. Just like real-world attacks, the Blue teams do not know when or how the Red team's attacks will occur or what methods may be used. Their job, whether it is a Red team attack or an actual assault, is to detect and respond to all security incidents. For this reason, the Blue teams are continuously on-call and must react to Red team breaches the same way they would for any other adversary.

Microsoft staffs separate full-time red teams and blue teams across Microsoft in various divisions that conduct operations both across services and within them. Referred to as *Red Teaming*, the approach is to test across Microsoft services' systems and operations using the same tactics, techniques, and procedures as real adversaries, against the live

production infrastructure, without the foreknowledge of the infrastructure and platform engineering or operations teams. This tests security detection and response capabilities, and helps identify production vulnerabilities, configuration errors, invalid assumptions, or other security issues in a controlled manner. Every Red team breach is followed by full disclosure between the Red team and Blue team, including service teams, to identify gaps, address findings and significantly improve breach response.

> **Note**
>
> No customer data is targeted during Red Teaming or live site penetration exercises. The tests are against Microsoft 365 and Azure infrastructure and platforms, as well as Microsoft's own tenants, applications, and data. Customer tenants, applications and data hosted in Azure, Dynamics 365, or Microsoft 365 are never targeted per the agreed to rules of engagement.

Joint exercises

At times, the Microsoft Blue and Red teams will conduct joint operations where the relationship during the operation is more partner than adversarial with a select set of employees from each team. These exercises are well coordinated between the teams to drive a more targeted set of outcomes through real-time collaboration between ethical hackers and responders. These 'purple team' exercises are highly tailored for each operation to maximize the opportunity, but fundamental to each operation is high-bandwidth information sharing and partnership to achieve the objectives.

Related articles

- Microsoft security incident management
- Microsoft security incident management: Detection and analysis
- Microsoft security incident management: Containment, eradication, and recovery
- Microsoft security incident management: Post-incident activity
- How to Log a Security Event Support Ticket
- Azure and Dynamics 365 breach notification under the GDPR

Microsoft security incident management: Detection and analysis

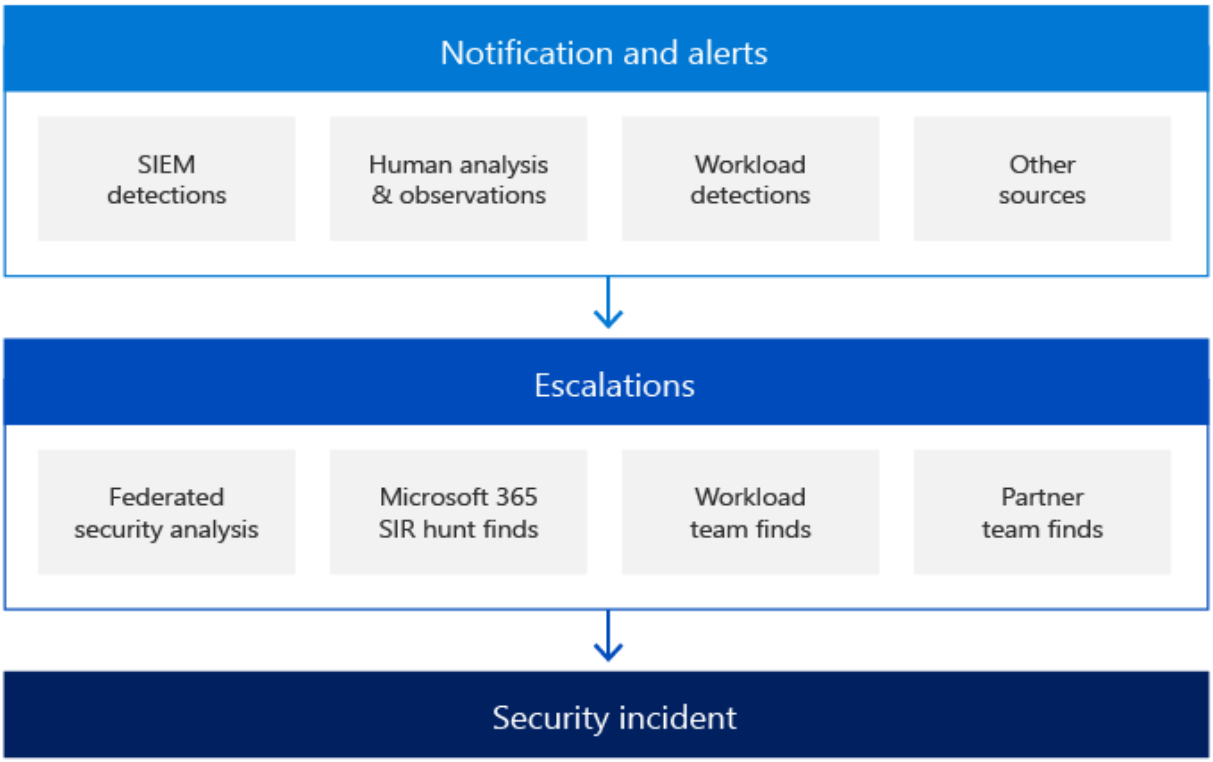
Article • 03/02/2023

To detect malicious activity, each of Microsoft's online services centrally logs security events and other data and perform various analytical techniques to find anomalous or suspicious activity. Log files are collected from Microsoft online services servers and infrastructure devices and stored in central and consolidated databases.

Microsoft takes a risk-based approach to detecting malicious activity. We use incident data and threat intelligence to define and prioritize our detections.

Employing a team of highly experienced, proficient, and skilled people is one of the most important pillars to success in the detection and analysis phase. Microsoft employs multiple service teams that include employees with competencies on all components within the stack, including the network, routers, firewalls, load balancers, operating systems, and applications.

The security detection mechanisms in Microsoft online services also include notifications and alerts that are initiated by different sources. Microsoft online services security response teams are the key orchestrators of the security incident escalation process. These teams receive all escalations and are responsible for analyzing and confirming the validity of the security incident.



One of the primary pillars of detection is notification:

- Each service team is responsible to log any action or event inside the service based on the requirements from the online service security team. All logs created by the different service teams are processed by a Security Information and Event Management (SIEM) solution with predefined security and detection rules. These rules evolve based on security team recommendations, on information learned from previous security incidents, to determine if there is any suspicious or malicious activity.
- If a customer determines that a security incident is underway, they may open a support case with Microsoft, which is assigned to the Microsoft communications team and turned into an escalation to all appropriate teams.

Azure, Dynamics 365, and Microsoft 365 service teams also use the intelligence gained in trend analysis through security monitoring and logging to detect abnormalities in Microsoft online services information systems that might indicate an attack or a security incident. Microsoft online services systems aggregate output from these logs in the production environment into centralized logging servers. From these centralized logging servers, logs are examined to spot trends throughout the production environment. Data aggregated in the centralized servers are securely transmitted into a logging service for advanced querying, dashboard building and detecting anomalous and malicious activity. The service also uses machine learning to detect anomalies with log output.

During the escalation phase and depending on the nature of the security incident, security response teams may engage one or more subject matter experts from various teams at Microsoft:

- Online Services Security and Compliance team
- Microsoft Threat Intelligence Center (MSTIC)
- Microsoft Security Response Center (MSRC)
- Corporate, External, and Legal Affairs (CELA)
- Azure Security
- Microsoft 365 engineering, and others.

Before an escalation to any security response team occurs, the service team is responsible for determining and setting the severity level of the security incident based on defined criteria such as:

- Privacy
- Impact
- Scope
- Number of affected tenants

- Region
- Service
- Details of the incident
- Specific customer industry or market regulations.

Incident prioritization is determined by using distinct factors, including but not limited to the functional impact of the incident, the informational impact of the incident, and the recoverability from the incident.

After receiving an escalation about a security incident, the security team organizes a virtual team (v-team) comprised of members from the Microsoft online service security response team, service teams, and the incident communication team. The v-team must then confirm the legitimacy of the security incident and eliminate any false positives. The accuracy of information provided by the indicators determined during the preparation phase is critical. By analyzing this information by category of vector attack, the v-team can determine if the security incident is a legitimate concern.

At the beginning of the investigation, the security incident response team records all information about the incident according to our case management policies. As the case progresses, we track ongoing actions and follow evidence handling standards for gathering, retaining, and securing this data throughout the incident lifecycle.

Some examples of these actions would include:

- A summary, which is a brief description of the incident and its potential impact
- The incident's severity and priority, which are derived by assessing the potential impact
- A list of all indicators identified which led to detection of the incident
- A list of any related incidents
- A list of all actions taken by the v-team
- Any gathered evidence, which will also be preserved for post-mortem analysis and future forensic investigations
- Recommended next steps and actions

After security incident confirmation, the primary goals of the security response team and the appropriate service team are to contain the attack, to protect the service(s) under attack, and to avoid a greater global impact. At the same time, the appropriate engineering teams work to determine the root cause and to prepare the first recovery plan.

In the next phase, the security response team identifies the customer(s) affected by the security incident, if any. The scope of effect can take some time to determine, based on region, datacenter, service, server farm, server, and so forth. The list of affected

customers is compiled by the service team and the corresponding Microsoft communications team, who then handle the customer notification process within contractual and compliance obligations.

Related articles

- Microsoft security incident management
- Microsoft security incident management: Preparation
- Microsoft security incident management: Containment, eradication, and recovery
- Microsoft security incident management: Post-incident activity
- How to Log a Security Event Support Ticket
- Azure and Dynamics 365 breach notification under the GDPR

Microsoft security incident management: Containment, eradication, and recovery

Article • 03/02/2023

Based on the analysis performed by the security response team, the service team, and others, an appropriate containment and recovery plan is developed to minimize the effect of the security incident. The appropriate service teams then apply that plan in production with support from the security response team.

Containment

After detecting a security incident, it is important to contain the intrusion before the adversary can access more resources or cause more damage. The primary goal of our security incident response procedures is to limit impact to customers or their data, or to Microsoft systems, services, and applications.

Eradication

Eradication is the process of eliminating the root cause of the security incident with a high degree of confidence. The goal is two-fold:

- to evict the adversary completely from the environment
- to mitigate the vulnerability (if known) that enabled or could enable the adversary to reenter the environment.

Depending on the nature of the incident, the scope of the security incident, the depth of the penetration and possible repercussions, the security response team will recommend that service teams adopt eradication techniques. Considering the potential business impact that may be caused by these eradication steps, these decisions will be made by service teams and the security response team after a detailed analysis and approval from the Executive Incident Manager (if necessary).

Recovery

As the response team gains a reasonable level of confidence that the adversary has been evicted from the environment and all known vulnerable paths have been eliminated, the individual service teams, will initiate restoration steps to bring the service

to a known and good configuration. These restoration steps will be in consultation with the security response team. This activity includes identifying the last known good state of the service, restoring from backups to this state, inspecting vulnerable attack paths in the restored state, etc. The security response team, in consultation with the service teams, will determine the best possible recovery plan for the environment.

A key aspect to the recovery is to have enhanced vigilance and controls in place to validate that the recovery plan has been successfully executed, and that no signs of breach exist in the environment.

Customer notification of security incident

If Microsoft determines that a security incident has occurred, we will notify you with undue delay, and within contractual and compliance requirements we have agreed to. After identifying all affected tenants, the corresponding communications team works to identify any relevant regulations that might apply to affected tenants. The communications team uses the appropriate communication channel defined in applicable regulations to notify the appropriate tenant contact.



Notification will include detailed information about the incident, such as a description of the incident, the effect on customer data, if any, actions taken by Microsoft, and/or suggested actions for customers to take to resolve the issue and prevent recurrence. Notification will be delivered to the designated administrator(s) of the Microsoft online services tenant. To ensure notifications are received, you should ensure that your administrators provide and maintain accurate contact information in their tenant profiles. In addition, depending on the nature of the incident, Microsoft 365 customers can also be notified via the Microsoft 365 Service Health Dashboard.

Related articles

- Microsoft security incident management
- Microsoft security incident management: Preparation
- Microsoft security incident management: Detection and analysis
- Microsoft security incident management: Post-incident activity
- How to Log a Security Event Support Ticket
- Azure and Dynamics 365 breach notification under the GDPR

Microsoft security incident management: Post-incident activity

Article • 03/02/2023

Postmortem

Some security incidents, especially those incidents that are customer impacting or result in a data breach, are subject to a full incident postmortem. The security response team conducts a detailed postmortem with all the parties involved in security incident response to:

- Document the sequence of events that caused the incident.
- Create a technical summary of the incident as supported by the evidence that includes the actors involved in the breach (if known). This summary will include how the response was executed and other key takeaways.
- Identify technical lapses, procedural failures, manual errors, process flaws, and communication glitches, and/or any previously unknown attack vectors that were identified during the security incident response.

The postmortem will directly influence Microsoft online service improvement, operational processes, and documentation by setting new priorities in the Microsoft online services engineering development cycle.

Documentation

All key technical findings in the postmortem process are captured in a report and service investments or fixes in the form of bugs or development change requests. These findings are followed-up with the appropriate engineering teams. For process failures and cross-organizational issues, issues are documented in the security response team's database and followed-up with the appropriate groups to address them.

Process improvement

Responding to a security incident in Microsoft online services involves coordination with multiple groups spread across different organizations within Microsoft, and potentially even appropriate external organizations such as law enforcement. We know that it is critical to evaluate our responses after every security incident for both sufficiency and completeness. For any identified improvements or changes, the security response team

evaluates the suggestions in consultation with the appropriate teams and stakeholders, and where appropriate incorporates them into standard operating procedures. All required changes, bugs, or service improvements identified during the security incident response or postmortem activity are logged and tracked in an internal Microsoft engineering database. All potential bugs or features are assigned to the appropriate owner. The Microsoft security response team reviews all entries until the issue is resolved.

Related articles

- Microsoft security incident management
- Microsoft security incident management: Preparation
- Microsoft security incident management: Detection and analysis
- Microsoft security incident management: Containment, eradication, and recovery
- How to Log a Security Event Support Ticket
- Azure and Dynamics 365 breach notification under the GDPR

Supplier management overview

Article • 05/25/2023

How does Microsoft manage risk related to suppliers?

Microsoft partners with third-party companies to help meet our customers' needs. These third-party companies are referred to as suppliers. Supplier security and privacy at Microsoft is governed by our Supplier Security and Privacy Assurance (SSPA) program, an enterprise-wide set of requirements for all suppliers who partner with Microsoft to deliver our online services. While the SSPA program provides comprehensive governance and management of our supplier base, individual business units may maintain additional requirements for their suppliers.

How does Microsoft's Supplier Security and Privacy Assurance (SSPA) Program protect customer data?

SSPA is a partnership between Microsoft Procurement, Corporate External and Legal Affairs, and Corporate Security to ensure suppliers adhere to Microsoft's privacy and security principles. The scope of SSPA covers all suppliers that process Personal Data or Microsoft Confidential Data. The SSPA program enrollment includes adherence to Microsoft's Data Protection Requirements (DPR). The DPR consist of security and privacy controls that suppliers must implement before beginning contracted work with Microsoft. All enrolled suppliers self-attest to compliance with the DPR annually.

DPR requirements are scoped based on six distinct data processing categories a supplier can be approved for as part of their enrollment in SSPA. These categories are used to identify the risk associated with the services a supplier provides to Microsoft. The supplier's data processing profile determines which DPR controls are considered in-scope to provide appropriate data protection. Suppliers who process data that is considered a higher risk must comply with all DPR requirements and may also need to provide independent verification of compliance. Microsoft purchasing tools validate the SSPA status of all suppliers, including compliance with applicable portions of the DPR, prior to allowing the procurement of that supplier.

What types of subprocessors provide services for Microsoft?

A 'subprocessor' is a third party that Microsoft engages whose duties include processing Microsoft Personal Data for which Microsoft is a processor. Microsoft's subprocessors fall into three categories. Each must demonstrate compliance with the SSPA before they can process customer data on Microsoft's behalf.

- **Technology** subprocessors that power technologies that are seamlessly integrated with Microsoft online services and in part power the Microsoft cloud functions. If a customer deploys one of these services, the subprocessors identified for that service may process, store, or otherwise access Customer Data or Personal Data while helping to provide that service.
- **Ancillary** subprocessors that provide services to help support, operate, and maintain online services. In such cases, the subprocessors identified may process, store, or otherwise access Customer Data and Personal Data (consisting of pseudonymized personal identifiers) while providing their ancillary services.
- **Contract Staff Organizations** provide contract staff who work side by side with Microsoft full-time employees to operate, deliver, and maintain Microsoft Online Services. In all such cases, Customer Data or Personal Data resides only on Microsoft systems, and is subject to Microsoft policies and supervision.

Additionally, **Microsoft data center infrastructure entities** provide the datacenter infrastructure on which the Microsoft Online Services run. The data within datacenters is encrypted, and no personnel within the datacenters can access it.

Technology and Ancillary third-parties are required to implement access controls in compliance with Microsoft's Data Protection Requirements (DPR). These requirements meet or exceed the contractual commitments Microsoft makes to its customers in the Product Terms. Suppliers who perform contract staff work are subject to the same access controls in place for Microsoft full-time employees.

How does Microsoft onboard suppliers?

Third-party suppliers are required to sign a Microsoft Master Agreement as part of the onboarding process. This agreement governs the relationship between Microsoft and its suppliers and ensures consistent management of supplier relationships. As part of onboarding, suppliers enroll in the SSPA and must complete all applicable requirements before they can be approved for any data processing categories. Microsoft business units are only able to create engagements with suppliers when the data processing

activity for the engagement matches data processing categories for which the supplier has been approved.

How does Microsoft notify customers of changes to suppliers who process their data?

Per the Microsoft Products and Services Data Protection Addendum (DPA), Microsoft makes additional commitments regarding notice periods for the addition of any subprocessor. Notice time frames depend on the type of data the subprocessor will process on behalf of Microsoft. As stated in the DPA, Microsoft commits to providing notice to customers at least six months in advance of any new subprocessor who will process Customer Data. For any other Personal Data, Microsoft will provide at least 30 days of notice. Notice is provided by the update of the Microsoft Online Services Subprocessor List.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the table below for validation of controls related to supplier management.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
ISO 27001/27002 Statement of Applicability Certificate	A.15.1: Information security in supplier relationships	April 24, 2023
ISO 27017 Statement of Applicability Certificate	A.15.1: Information security in supplier relationships	April 24, 2023
ISO 27018 Statement of	A.8.1: Disclosure of subcontracted PII processing	April 24, 2023

External audits	Section	Latest report date
Applicability Certificate		
SOC 2 SOC 3	SOC2-25: Supplier risk management C5-2: Supplier risk profile review	November 17, 2023

Microsoft 365

Expand table

External audits	Section	Latest report date
FedRAMP (Office 365)	CA-3: System interconnections IA-4: Identifier management PS-6: Access agreements PS-7: Third-party personnel security SA-4: Acquisitions process SA-9: External information system services SA-12: Supply chain protection	July 31, 2023
ISO 27001/27002/27017 Statement of Applicability Certification (27001/27002) Certification (27017)	A.15.1: Information security in supplier relationships	March 2023
ISO 27018 Statement of Applicability Certificate	A.8.1: Disclosure of subcontracted PII processing	March 2023
SOC 2	CA-53: Third-party monitoring	January 3, 2023

Resources

- Supplier Security Privacy and Assurance Program

Feedback

Was this page helpful? Yes No

Supplier Security and Privacy Assurance (SSPA) program

Article • 04/19/2023

> **Important**
>
> The information presented in this article is on behalf of the Supplier Security and Privacy Assurance (SSPA) team. The most up to date information is available here. If there is a conflict between the information presented in this article and the SSPA page, the SSPA page will supersede the information in this article.

Microsoft believes that privacy is a fundamental right. In the mission to empower every individual and organization on the planet to achieve more, Microsoft strives to earn and maintain the trust of their customers.

Strong privacy and security practices are critical to this mission, essential to trust, and in several jurisdictions required by law. The standards captured in Microsoft's privacy and security policies reflect our values as a company and extend to suppliers that process Personal and Confidential Data on our behalf.

The Supplier Security and Privacy Assurance (SSPA) Program delivers Microsoft's baseline data processing instructions to suppliers in the form of the Microsoft Supplier Data Protection Requirements (DPR).

> **Note**
>
> Suppliers may have to meet additional organizational level requirements that are decided and communicated outside of SSPA by the Microsoft group responsible for the engagement with the supplier.

SSPA program overview

SSPA is a partnership between Microsoft Procurement, Corporate External and Legal Affairs, and Corporate Security to ensure privacy and security principles are followed by suppliers. The scope of SSPA covers all suppliers globally that process Personal Data and/or Microsoft Confidential Data.

SSPA enables the supplier to make Data Processing Profile selections that align with the goods and/or services suppliers are contracted to perform. These selections trigger corresponding requirements to provide compliance assurances.

All enrolled suppliers must complete an annual self-attestation of DPR compliance. A supplier's Data Processing Profile determines whether the full DPR is issued or if a subset of requirements applies. Suppliers that process data that Microsoft considers higher risk may also need to meet additional requirements, such as providing independent verification of compliance. Suppliers that are on a published Microsoft subprocessor list will also be asked to provide independent verification of compliance.

> **ⓘ Important**
>
> Compliance activities determine an SSPA status of Green (compliant) or Red (noncompliant). Microsoft purchasing tools validate the SSPA status is Green (for each supplier in scope for SSPA) prior to allowing an engagement to move forward.

SSPA scope

To help determine whether the supplier processes Personal Data and/or Microsoft Confidential Data, see the list of examples in the following tables. These are examples and not an exhaustive list.

Personal data by data type

Examples include but are not limited to:

Data Type	Examples

Data Type	Examples
Sensitive data	• Data related to children • Genetic data, biometric data, or health data • Racial or ethnic origin • Political, religious, or philosophical beliefs, opinions, and affiliations • Trade union membership • A natural person's sex life or sexual orientation • Immigration status (visa, work authorization, etc.) • Government identifiers (passport, driver's license, visa, social security numbers, national identity numbers) • Precise user location data (within 300 meters) • Personal bank account numbers • Credit card number and expiration date
Customer content data	• Documents, photos, videos, music, etc. • Browsing history, interests, and favorites linking, typing, and speech utterance (voice/audio and/or chat/bot)
Captured and generated data	• Imprecise location data • IP address • Device preferences and personalization service usage for websites, webpage click tracking • Contact data such as name, address, phone number, email address, date of birth, dependent, and emergency contacts • Fraud and risk assessment, background check • Metadata and telemetry
Account data	• Payment instrument data • Credit card number and expiration date • Bank routing information • Bank account number • Credit requests or line of credit • Tax documents and identifiers • Investment or expense data
End-user Pseudonymized Information (EUPI) *Identifiers created by Microsoft to identify users of Microsoft products and services*	• Globally Unique Identifier (GUID) • Unique Identifier (PUID) • Hashed End-User Identifiable Information (EUII) • Session IDs • Device IDs • Diagnostic data log data

Data Type	Examples
Online customer data	• Microsoft online enterprise customer (Azure tenant, M365 tenant) • Microsoft enterprise customer (on-premises customer) • Account data (billing data, e-commerce) • Survey/event registration/training

Microsoft confidential data by data class

Examples include but are not limited to:

Data Class	Examples
Highly Confidential	• Information concerning or related to the development, testing, or manufacturing of Microsoft products or components of Microsoft products. Microsoft software, online services, or hardware sold commercially in any channel are considered "Microsoft Product" • Microsoft device pre-release marketing information • Unannounced Microsoft corporate financial data subject to SEC rules
Confidential	• Microsoft product license keys on behalf of Microsoft for distribution via any method • Information concerning or related to the development or testing of Microsoft internal Line of Business (LOB) applications • Microsoft pre-release marketing materials for Microsoft software and services such as Office, SQL, Azure, etc. • Written, design, electronic, or print documentation for any Microsoft services or products, such as devices (process or procedure guides, configuration data, etc.)

Data Processing Profile

Microsoft suppliers have control over their SSPA Data Processing Profile, allowing suppliers to decide which engagements they want to be eligible to perform.

Microsoft business groups are only able to create engagements with suppliers where the data processing activity matches the approvals the supplier has obtained.

Suppliers are able to update their Data Processing Profile at any time during the year if there are no open tasks. When a change is made, the corresponding activity is issued

and must be completed before the approvals are secured. The existing, completed approvals apply until newly issued requirements are completed.

If the newly executed tasks are not completed within the 90-day time frame allowed, the SSPA status is updated to Red (non-compliant), and the account is deactivated from Microsoft Accounts Payable systems.

Data processing scope considerations

- **Confidential:** If the supplier only processes Microsoft Confidential Data, their profile will show a Confidential flag. The supplier will not be able to process any Personal Data.
- **Personal**, **Confidential:** If the supplier processes Personal Data, their profile will indicate the Personal, Confidential flag.
- **Processing Location At Microsoft or Customer:** If the supplier will only process data within Microsoft systems, using Microsoft credentials and are subject to Microsoft security and privacy policies, this option will be selected on the supplier's profile.
- **At Supplier:** If the condition "At Microsoft or Customer" (as described previously) does not apply, this is the option that will be reflected in the supplier's profile.

Data processing role considerations

- **Controller:** (covers independent and joint controllers) If the supplier is both a controller and a processor (for different engagements), the supplier will select "Processor".
- **Processor** When the supplier processes data on behalf of Microsoft.
- **Subprocessor:** A subprocessor is a third party that Microsoft engages to perform, where their performance includes processing of Microsoft Personal Data for which Microsoft is a Processor. Suppliers cannot self-identify as a subprocessor at Microsoft because it requires pre-approval by internal Privacy teams. Suppliers can only be a subprocessor when Microsoft is the Data Processor and the supplier processes qualifying Enterprise Personal Datatypes. Subprocessors will have additional contract and compliance requirements, including a Data Protection Addendum and an Independent Assessment, and additional certification requirements.

- **Payment Card Processing:** If any part of the data processed by a supplier includes data to support credit card or other payment card processing on behalf of Microsoft. This approval allows a supplier to engage in payment card processing engagements.
- **Software:** Microsoft Procurement directs buyers through an intake process for all software purchases, this includes various checks including SSPA triage to decide if the supplier providing the software is in-scope for SSPA management. If SSPA is required, suppliers may also need to identify that the 'Software as a Service' (SaaS) profile choice applies. For SSPA enrolled suppliers, this can be done when completing the Data Processing Profile in the Microsoft Supplier Compliance Portal. For SSPA compliance purposes, view SaaS broadly to also include platform as a service (PaaS), and infrastructure as a service (IaaS).
- **Software as a Service (SaaS):** SaaS allows users to connect to and use cloud-based applications over the Internet. Microsoft defines SaaS as software based on common code used in a one-to-many model on a pay-for-use basis or as a subscription based on use metrics. The cloud service provider develops and maintains cloud-based software, provides automatic software updates, and makes software available to its customers via the internet on a one-to-many, pay-as-you-go basis. This method of software delivery and licensing allows software to be accessed online via a subscription rather than bought and installed on each individual computer.

> ⓘ **Note**
>
> Most SaaS suppliers will need to add the Subcontractor approval in the Microsoft Supplier Compliance Portal if the Personal Data or Microsoft Confidential data is hosted on a 3rd party platform.

- **Use of Subcontractors:** This flag is required if the supplier uses subcontractors to perform any part of the contracted work. This also includes Freelancers.

Assurance requirements

The approvals selected in the supplier's Data Processing Profile assists SSPA in assessing the risk level across the supplier's engagement(s). SSPA compliance requirements differ based on the Data Processing Profile and associated approvals.

There are also combinations that may elevate or reduce compliance requirements. The combinations are captured in the Requirements based on profile approvals section.

If the supplier's profile includes Software as a Service (SaaS), subcontractors, website hosting, or payment cards, additional assurances are required.

Self-attestation to the DPR

All suppliers enrolled in SSPA must complete a self-attestation of compliance to the DPR within 90 days of receiving the request. This request must be provided on an annual basis but may be more frequent if the Data Processing Profile is updated mid-year. Supplier accounts will change to an SSPA status of Red (non-compliant) if the 90-day period is exceeded. New in-scope purchase orders cannot process until the SSPA status turns to Green (compliant).

Newly enrolled suppliers must complete issued requirements to secure a SSPA status of Green (compliant) before engagements can begin.

Applicability

Suppliers are expected to respond to all applicable DPR requirements issued per the Data Processing Profile. It is expected that, within the issued requirements, a few may not apply to the goods or services the supplier provides to Microsoft. These can be marked as 'does not apply' with a detailed comment for SSPA reviewers to validate.

DPR submissions are reviewed by the SSPA team for any selections of 'does not apply', 'local legal conflict', or 'contractual conflict' against issued requirements.

Independent assessment requirement

Please see the Requirements by approvals section to see the data processing approvals that trigger this requirement.

Suppliers have the option to change approvals by updating their Data Processing Profile. However, if the supplier has a Data Processing Role of "Subprocessor", the supplier cannot change this approval and will be required to have an independent assessment conducted annually.

The Requirements based on profile approvals section includes acceptable certification alternatives if you elect not to use an independent assessor to verify compliance to the DPR (when applicable, such as for SaaS suppliers, website hosting suppliers or suppliers with Subcontractors). The ISO 27701 (privacy) and ISO 27001 (security) are relied on as providing close mapping to the DPR.

If a supplier is a healthcare provider in the United States or covered entity, Microsoft will accept a HITRUST report for privacy and security coverage.

SSPA may execute an independent assessment manually if circumstances beyond standard triggers warrant additional due diligence. Examples include a request from division privacy or security; validation of data incident remediation; or requirement for automated data subject rights execution.

PCI DSS certification requirement

The Payment Card Industry Data Security Standard (PCI DSS) is a framework for developing robust payment card data security that includes prevention, detection, and appropriate reaction to security incidents. The framework was developed by the PCI Security Standards Council, a self-regulatory industry organization. The purpose of the PCI DSS requirements is to identify technology and process vulnerabilities that pose risks to the security of cardholder data that is processed.

Microsoft is required to comply with these standards. If a supplier handles payment card information on Microsoft's behalf, we require evidence of adherence to these standards.

Depending on the volume of transactions processed, a supplier will either be required to have a Qualified Security Assessor certify compliance or can complete a self-assessment questionnaire form.

Payment card brands set the thresholds for assessment type, typically:

- Level 1: Provide a Third Party Assessor PCI AOC certificate
- Level 2 or 3: Provide a PCI DSS Self-Assessment Questionnaire (SAQ) signed by the supplier's officer.

Software as a Service requirement

Suppliers that met the SaaS definition included on the Data Processing Profile may be required to provide a valid ISO 27001 certification.

Use of subcontractors

Microsoft considers the use of subcontractors a high-risk factor. Suppliers using subcontractors who will process Personal and or Microsoft Confidential Data must disclose those subcontractors. Additionally, the supplier should also disclose the countries where that personal data will be processed by each subcontractor.

Requirements based on profile approvals

#	Profile	Assurance Requirements	Independent Assurance Options
1	**Scope:** Personal, Confidential **Processing Location:** At Microsoft or Customer Processing **Role:** Processor or Controller **Data Class:** Confidential or Highly Confidential **Payment Cards:** Not Applicable **SaaS:** Not Applicable **Use of Subcontractors:** Not Applicable **Website Hosting:** Not Applicable	Self-attestation of compliance to the DPR	
2	**Scope:** Confidential **Processing Location:** At Supplier **Role:** N/A **Data Class:** Confidential **Payment Cards:** Not Applicable **SaaS:** Not Applicable **Use of Subcontractors:** Not Applicable **Website Hosting:** Not Applicable	Self-attestation of compliance to the DPR	
3	**Scope:** Confidential Processing **Processing Location:** At Supplier **Role:** Processor **Data Class:** Highly Confidential **Payment Cards:** Not Applicable **SaaS:** Not Applicable **Use of Subcontractors:** Not Applicable **Website Hosting:** Not Applicable	Self-attestation of compliance to the DPR **and** Independent Assurance of compliance	**Independent Assurance options:** 1. Complete an Independent Assessment against the DPR 2. Submit ISO 27001

#	Profile	Assurance Requirements	Independent Assurance Options
4	**Scope:** Personal, Confidential **Processing Location:** At Supplier Processing **Role:** Processor **Data Class:** Highly Confidential **Payment Cards:** Not Applicable **SaaS:** Not Applicable **Use of Subcontractors:** Not Applicable **Website Hosting:** Not Applicable	Self-attestation of compliance to the DPR **and** Independent Assurance of compliance	**Independent Assurance options:** 1. Complete an Independent Assessment against the DPR 2. Independent Assessment against sections A-I of the DPR and ISO 27001 3. Submit ISO 27701 and ISO 27001
5	**Scope:** Personal, Confidential **Processing Location:** At Supplier **Role:** Processor **Data Class:** Confidential **Payment Cards:** Not Applicable **SaaS:** Not Applicable **Use of Subcontractors:** Not Applicable **Website Hosting:** Not Applicable	Self-attestation of compliance to the DPR	
6	**Scope:** Personal, Confidential **Processing Location:** At Supplier **Role:** Controller **Data Class:** Confidential or Highly Confidential **Payment Cards:** Not Applicable **SaaS:** Not Applicable **Use of Subcontractors:** Not Applicable **Website Hosting:** Not Applicable	Self-attestation of compliance to the DPR	

#	Profile	Assurance Requirements	Independent Assurance Options
7	**Scope:** Personal, Confidential **Processing Location:** Any **Role:** Subprocessor (This role is determined by Microsoft – profile will read "Subprocessor Approval: Yes") **Data Class:** Confidential or Highly Confidential **Payment Cards:** Not Applicable **SaaS:** Not Applicable **Use of Subcontractors:** Not Applicable **Website Hosting:** Not Applicable	Self-attestation of compliance to the DPR **and** Independent Assurance of compliance	**Independent Assurance options:** 1. Complete an Independent Assessment against the DPR 2. Independent Assessment against sections A-I of the DPR and ISO 27001 3. Submit ISO 27701 and ISO 27001

Threat and vulnerability management overview

Article • 09/12/2023

How do Microsoft online services conduct vulnerability management?

No matter how well a system is designed, its security posture can degrade over time. Machines can go unpatched, inadvertent configuration changes can be introduced, and regressions to security code can accumulate. All these issues can make a system less secure than when it was initially deployed. Microsoft has built automation to continually assess our systems for this kind of degradation, enabling us to act immediately to correct issues in our security posture.

Microsoft online services use machine state scanning to make sure the machines comprising our infrastructure are up to date with the latest patches and that their base configurations correctly align with relevant frameworks. Machine state scanning uses patching, anti-malware, vulnerability scanning, and configuration scanning (PAVC). Microsoft online services apply effective PAVC by installing a custom security agent on each asset during deployment. This security agent enables machine state scanning and reports results to our service teams.

How do Microsoft online services ensure service infrastructure is up to date with the latest security patches?

Patch management mitigates vulnerabilities by ensuring Microsoft online services systems are updated quickly when new security patches are released. Microsoft prioritizes new security patches and other security updates according to risk. Microsoft online service security teams analyze available security patches to determine their risk level in the context of our production environments. Their analysis includes severity scores based on the Common Vulnerability Scoring System (CVSS) along with other risk factors.

Microsoft service teams review the analysis from the security team and update their service components and baseline images with applicable patches within the appropriate remediation timeframe. Security patches are subject to the change management process

to ensure adequate testing and management approval before deployment to production environments. Deployment of security patches occurs in stages to enable rollback if a security patch causes unexpected issues.

Service teams use vulnerability scan results to validate security patch deployment on applicable system components. Any overdue vulnerabilities are reported daily and reviewed by management monthly to measure the breadth and depth of patch coverage across the environment and hold ourselves accountable for timely patching.

How does Microsoft conduct vulnerability and configuration scanning?

Microsoft's security agent is installed during asset deployment and enables fully automated vulnerability and configuration scanning. The security agent uses industry-standard tools to detect known vulnerabilities and security misconfigurations. Production assets are scheduled for daily, automatic scans with the most recent vulnerability signatures. The results of these scans are collected in a secure, central storage service, and automated reporting makes results available to service teams.

Service teams review scan results using dashboards that report aggregate scan results to provide comprehensive reporting and trend analysis. Vulnerabilities detected in scans are tracked in these reports until they're remediated. When vulnerability scans indicate missing patches, security misconfigurations, or other vulnerabilities in the environment, service teams use these reports to target the affected components for remediation. Vulnerabilities discovered through scanning are prioritized for remediation based on their Common Vulnerability Scoring System (CVSS) scores and other relevant risk factors.

How does Microsoft defend against malware?

Microsoft uses comprehensive anti-malware software to protect Microsoft online services against viruses and other malware. Baseline operating system images used by Microsoft online services include this software to maximize coverage throughout the environment.

Every endpoint in Microsoft online services performs a full anti-malware scan at least weekly. Additional real-time scans are performed on all files as they're downloaded, opened, or executed. These scans use known malware signatures to detect malware and prevent malware execution. Microsoft's anti-malware software is configured to download the most recent malware signatures daily to ensure scans are conducted with the most up-to-date information. In addition to signature-based scans, Microsoft anti-

malware software uses pattern-based recognition to detect and prevent suspicious or anomalous program behavior.

When our anti-malware products detect viruses or other malware, they automatically generate an alert for Microsoft security response teams. In many cases, our anti-malware software can prevent the execution of viruses and other malware in real time without human intervention. When this prevention isn't possible, Microsoft security response teams resolve malware incidents using the security incident response process.

How does Microsoft detect new or unreported vulnerabilities?

Microsoft supplements automated scanning with sophisticated machine learning to help detect suspicious activity that might indicate the presence of unknown vulnerabilities. Regular penetration testing by internal Microsoft teams and independent auditors provides an additional mechanism for discovering and remediating vulnerabilities before they can be exploited by real-world attackers. Microsoft conducts internal penetration testing using "Red Teams" of Microsoft ethical hackers. Customer systems and data are never the targets of penetration testing, but lessons learned from penetration tests help Microsoft validate its security controls and defend against new kinds of attacks. Microsoft also uses bug bounty programs to incentivize disclosure of new vulnerabilities, enabling them to be mitigated as soon as possible.

Related external regulations & certifications

Microsoft's online services are regularly audited for compliance with external regulations and certifications. Refer to the following table for validation of controls related to threat and vulnerability management.

Azure and Dynamics 365

Expand table

External audits	Section	Latest report date
ISO 27001/27002 Statement of Applicability Certificate	A.12.6.1: Management of technical vulnerabilities	April 24, 2023

External audits	Section	Latest report date
ISO 27017 Statement of Applicability Certificate	A.12.6.1: Management of technical vulnerabilities	April 24, 2023
ISO 27018 Statement of Applicability Certificate	A.12.6.1: Management of technical vulnerabilities	April 24, 2023
SOC 1 SOC 2 SOC 3	VM-3: Anti-malware monitoring VM-5: Patching VM-6: Vulnerability Scanning	November 17, 2023

Microsoft 365

Expand table

External audits	Section	Latest report date
FedRAMP (Office 365)	CA-7: Continuous monitoring CA-8: Penetration testing RA-3: Risk assessment RA-5: Vulnerability scanning SI-2: Flaw remediation SI-5: Security alerts, advisories, and directives	July 31, 2023
ISO 27001/27002/27017 Statement of Applicability Certification (27001/27002) Certification (27017)	A.12.6.1: Management of technical vulnerabilities	March 2023
SOC 1	CA-27: Vulnerability scanning	January 3, 2023
SOC 2	CA-24: Internal risk assessment CA-27: Vulnerability scanning	January 3, 2023

Resources

- Video: Microsoft 365 vulnerability management

- Azure Commercial Penetration Test Report: 2022
- Microsoft Office 365 Third-Party Vulnerability Assessment: 2022
- Additional penetration test and vulnerability assessments

Feedback

Was this page helpful? Yes No

Attack simulation in Microsoft 365

Article • 03/02/2023

Microsoft continuously monitors and explicitly tests for weaknesses and vulnerabilities in tenant boundaries, including monitoring for intrusion, permission violation attempts, and resource starvation. We also use multiple internal systems to continuously monitor for inappropriate resource utilization, which if detected, triggers built-in throttling.

Microsoft 365 has internal monitoring systems that continuously monitor for any failure and drive automated recovery when failure is detected. Microsoft 365 systems analyze deviations in service behavior and initiate self-healing processes that are built into the system. Microsoft 365 also uses outside-in monitoring in which monitoring is performed from multiple locations both from trusted third-party services (for independent SLA verification) and our own datacenters to raise alerts. For diagnostics, we have extensive logging, auditing, and tracing. Granular tracing and monitoring help us isolate issues and perform fast and effective root cause analysis.

While Microsoft 365 has automated recovery actions where possible, Microsoft on-call engineers are available 24x7 to investigate all Severity 1 security escalations, and post-mortem reviews of every service incident contributes to continuous learning and improvement. This team includes support engineers, product developers, program managers, product managers, and senior leadership. Our on-call professionals provide timely backup and often can automate recovery actions, so that next time an event occurs, it can be self-healed.

Microsoft performs a thorough post-incident review each time a Microsoft 365 security incident occurs regardless of the magnitude of impact. A post-incident review consists of an analysis of what happened, how we responded and how we prevent similar incidents in the future. In the interest of transparency and accountability, we share post-incident reviews for any major service incidents with affected customers. For specific details, see Microsoft security incident management.

Assume breach methodology

Based on detailed analysis of security trends, Microsoft advocates and highlights the need for other investments in reactive security processes and technologies that focus on detection and response to emerging threats, rather than solely the prevention of those threats. Because of changes in the threat landscape and in-depth analysis, Microsoft refined its security strategy beyond just preventing security breaches to one better

equipped to deal with breaches when they do occur; a strategy that considers major security events not as a matter of if, but when.

While Microsoft's assume breach practices have been in place for many years, many customers are unaware of the work being done behind the scenes to harden the Microsoft cloud. Assume breach is a mindset that guides security investments, design decisions, and operational security practices. Assume breach limits the trust placed in applications, services, identities, and networks by treating them all—internal and external—as insecure and already compromised. Although the assume Breach strategy was not borne from an actual breach of any Microsoft enterprise or cloud services, it was recognized that many organizations, across the industry, were being breached despite all attempts to prevent it. While preventing breaches is a critical part of any organization's operations, those practices must be continuously tested and augmented to effectively address modern adversaries and advanced persistent threats. For any organization to prepare for a breach, they must first build and maintain robust, repeatable, and thoroughly tested security response procedures.

While prevent breach security processes, such as threat modeling, code reviews, and security testing are very useful as part of the Security Development Lifecycle, assume breach provides numerous advantages that help account for overall security by exercising and measuring reactive capabilities in the event of a breach.

At Microsoft, we set out to accomplish this through ongoing war-games exercises and live site penetration testing of our security response plans with the goal of improving our detection and response capability. Microsoft regularly simulates real-world breaches, conducts continuous security monitoring, and practices security incident management to validate and improve the security of Microsoft 365, Azure, and other Microsoft cloud services.

Microsoft executes its assume breach security strategy using two core groups:

- Red Teams (attackers)
- Blue Teams (defenders)

Both Microsoft Azure and Microsoft 365 staff separate full-time Red Teams and Blue Teams.

Referred to as "Red Teaming", the approach is to test Azure and Microsoft 365 systems and operations using the same tactics, techniques and procedures as real adversaries, against the live production infrastructure, without the foreknowledge of the Engineering or Operations teams. This tests Microsoft's security detection and response capabilities, and helps identify production vulnerabilities, configuration errors, invalid assumptions, and other security issues in a controlled manner. Every Red Team breach is

followed by full disclosure between both teams to identify gaps, address findings, and improve breach response.

NOTE: No customer data is deliberately targeted during Red Teaming or live site penetration testing. The tests are against Microsoft 365 and Azure infrastructure and platforms, as well as Microsoft's own tenants, applications, and data. Customer tenants, applications, and content hosted in Microsoft 365 or Azure are never targeted.

Red Teams

The Red Team is a group of full-time staff within Microsoft that focuses on breaching Microsoft's infrastructure, platform, and Microsoft's own tenants and applications. They are the dedicated adversary (a group of ethical hackers) performing targeted and persistent attacks against Online Services (Microsoft infrastructure, platforms, and applications but not end-customers' applications or content).

The role of the Red Team is to attack and penetrate environments using the same steps as an adversary:



Among other functions, red teams specifically attempt to breach tenant isolation boundaries to find bugs or gaps in our isolation design.

To help scale attack simulation efforts, the Red Team has created an automated attack emulation tool that runs safely in specific Microsoft 365 environments on a recurring basis. The tool has a wide variety of predefined attacks that are constantly expanded and improved to help reflect the evolving threat landscape. In addition to broadening the coverage of Red Team testing, it helps the Blue Team validate and improve their security monitoring logic. Regular, ongoing attack emulation provides the Blue Team with a consistent and diverse stream of signals that are compared and validated against expected responses. This helps lead to improvements in Microsoft 365's security monitoring logic and response capabilities.

Blue Teams

The Blue Team is composed of either a dedicated set of security responders or members from across the security incident response, Engineering, and Operations organizations. Regardless of their make-up, they are independent and operate separately from the Red Team. The Blue Team follows established security processes and uses the latest tools and

technologies to detect and respond to attacks and penetration. Just like real-world attacks, the Blue Team does not know when or how the Red Team's attacks will occur or what methods may be used. Their job, whether it is a Red Team attack or an actual assault, is to detect and respond to all security incidents. For this reason, the Blue Team is continuously on-call and must react to Red Team breaches the same way they would for any other breach.

When an adversary, such as a Red Team, has breached an environment, the Blue Team must:

- Gather evidence left by the adversary
- Detect the evidence as an indication of compromise
- Alert the appropriate Engineering and Operation team(s)
- Triage the alerts to determine whether they warrant further investigation
- Gather context from the environment to scope the breach
- Form a remediation plan to contain or evict the adversary
- Execute the remediation plan and recover from breach

These steps form the security incident response that runs parallel to the adversary's, as shown below:



Red Team breaches allow for exercising the Blue Team's ability to detect and respond to real-world attacks end-to-end. Most importantly, it allows for practiced security incident response prior to a genuine breach. Additionally, because of Red Team breaches, the Blue Team enhances their situational awareness, which can be valuable when dealing with future breaches (whether from the Red Team or another adversary). Throughout the detection and response process, the Blue Team produces actionable intelligence and gains visibility into the actual conditions of the environment(s) they are trying to defend. Frequently this is accomplished via data analysis and forensics, performed by the Blue Team, when responding to Red Team attacks and by establishing threat indicators, such as indicators of compromise. Much like how the Red Team identifies gaps in the security story, blue teams identify gaps in their ability to detect and respond. Furthermore, since the red teams model real-world attacks, the Blue Team can be accurately assessed on their ability, or inability, to deal with determined and persistent adversaries. Finally, Red Team breaches measure both readiness and impact of our breach response.

Malware and ransomware protection in Microsoft 365

Article • 03/02/2023

Protecting customer data from malware

Malware consists of viruses, spyware and other malicious software. Microsoft 365 includes protection mechanisms to prevent malware from being introduced into Microsoft 365 by a client or by a Microsoft 365 server. The use of anti-malware software is a principal mechanism for protection of Microsoft 365 assets from malicious software. The anti-malware software detects and prevents computer viruses, malware, rootkits, worms, and other malicious software from being introduced into any service systems. Anti-malware software provides both preventive and detective control over malicious software.

Each anti-malware solution in place tracks the version of the software and what signatures are running. The automatic download and application of signature updates at least daily from the vendor's virus definition site is centrally managed by the appropriate anti-malware tool for each service team. The following functions are centrally managed by the appropriate anti-malware tool on each endpoint for each service team:

- Automatic scans of the environment
- Periodic scans of the file system (at least weekly)
- Real-time scans of files as they're downloaded, opened, or executed
- Automatic download and application of signature updates at least daily from the vendor's virus definition site
- Alerting, cleaning, and mitigation of detected malware

When anti-malware tools detect malware, they block the malware and generate an alert to Microsoft 365 service team personnel, Microsoft 365 Security, and/or the security and compliance team of the Microsoft organization that operates our datacenters. The receiving personnel initiate the incident response process. Incidents are tracked and resolved, and post-mortem analysis is performed.

Exchange Online Protection against malware

All email messages for Exchange Online travel through Exchange Online Protection (EOP), which quarantines and scans in real time all email and email attachments both entering and leaving the system for viruses and other malware. Administrators don't

need to set up or maintain the filtering technologies; they're enabled by default. However, administrators can make company-specific filtering customizations using the Exchange admin center.

Using multiple anti-malware engines, EOP offers multilayered protection that's designed to catch all known malware. Messages transported through the service are scanned for malware (including viruses and spyware). If malware is detected, the message is deleted. Notifications may also be sent to senders or administrators when an infected message is deleted and not delivered. You can also choose to replace infected attachments with either default or custom messages that notify the recipients of the malware detection.

The following helps provide anti-malware protection:

- **Layered Defenses Against Malware** - Multiple anti-malware scan engines used in EOP help protect against both known and unknown threats. These engines include powerful heuristic detection to provide protection even during the early stages of a malware outbreak. This multi-engine approach has been shown to provide significantly more protection than using just one anti-malware engine.
- **Real-time Threat Response** - During some outbreaks, the anti-malware team may have enough information about a virus or other form of malware to write sophisticated policy rules that detect the threat even before a definition is available from any of the engines used by the service. These rules are published to the global network every 2 hours to provide your organization with an extra layer of protection against attacks.
- **Fast Anti-Malware Definition Deployment** - The anti-malware team maintains close relationships with partners who develop anti-malware engines. As a result, the service can receive and integrate malware definitions and patches before they're publicly released. Our connection with these partners often allows us to develop our own remedies as well. The service checks for updated definitions for all anti-malware engines every hour.

Microsoft Defender for Office 365

Microsoft Defender for Office 365 is an email filtering service that provides additional protection against specific types of advanced threats, including malware and viruses. Exchange Online Protection currently uses a robust and layered anti-virus protection powered by multiple engines against known malware and viruses. Microsoft Defender for Office 365 extends this protection through a feature called Safe Attachments, which protects against unknown malware and viruses, and provides better zero-day protection to safeguard your messaging system. All messages and attachments that don't have a known virus/malware signature are routed to a special hypervisor environment, where a

behavior analysis is performed using a variety of machine learning and analysis techniques to detect malicious intent. If no suspicious activity is detected, the message is released for delivery to the mailbox.

Exchange Online Protection also scans each message in transit in Microsoft 365 and provides time of delivery protection, blocking any malicious hyperlinks in a message. Attackers sometimes try to hide malicious URLs with seemingly safe links that are redirected to unsafe sites by a forwarding service after the message has been received. Safe Links proactively protects your users if they select such a link. That protection remains every time they select the link, and malicious links are dynamically blocked while good links are accessible.

Microsoft Defender for Office 365 also offers rich reporting and tracking capabilities, so you can gain critical insights into who is getting targeted in your organization and the category of attacks you're facing. Reporting and message tracing allows you to investigate messages that have been blocked due to an unknown virus or malware, while the URL trace capability allows you to track individual malicious links in the messages that have been clicked.

For more information about Microsoft Defender for Office 365, see Exchange Online Protection and Microsoft Defender for Office 365.

SharePoint Online and OneDrive for Business Protection Against Ransomware

There are many forms of ransomware attacks, but one of the most common forms is where a malicious individual encrypts a user's important files and then demands something from the user, such as money or information, in exchange for the key to decrypt them. Ransomware attacks are on the rise, particularly those that encrypt files that are stored in the user's cloud storage. For more information about ransomware, see the Microsoft Defender Security Intelligence site.

Versioning helps to protect SharePoint Online lists and SharePoint Online and OneDrive for Business libraries from some, but not all, of these types of ransomware attacks. Versioning is enabled by default in OneDrive for Business and SharePoint Online. Since versioning is enabled in SharePoint Online site lists, you can look at earlier versions and recover them, if necessary. That enables you to recover versions of items that pre-date their encryption by the ransomware. Some organizations also retain multiple versions of items in their lists for legal reasons or audit purposes.

SharePoint Online and OneDrive for Business Recycle Bins

SharePoint Online administrators can restore a deleted site collection by using the SharePoint Online admin center. SharePoint Online users have a Recycle Bin where deleted content is stored. They can access the Recycle Bin to recover deleted documents and lists, if they need to. Items in the Recycle Bin are retained for 93 days. The following data types are captured by the Recycle Bin:

- Site collections
- Sites
- Lists
- Libraries
- Folders
- List items
- Documents
- Web Part pages

Site customizations made through SharePoint Designer aren't captured by the Recycle Bin. For more information, see Restore deleted items from the site collection recycle bin. See also, Restore a deleted site collection.

Versioning doesn't protect against ransomware attacks that copy files, encrypt them, and then delete the original files. However, end-users can leverage the Recycle Bin to recover OneDrive for Business files after a ransomware attack occurs.

The following section goes into more detail on the defenses and controls Microsoft uses to mitigate the risk of cyberattack against your organization and its assets.

How Microsoft mitigates risks from a ransomware attack

Microsoft has built in defenses and controls it uses to mitigate the risks of a ransomware attack against your organization and its assets. Assets can be organized by domain with each domain having its own set of risk mitigations.

Domain 1: Tenant level controls

The first domain is the people that make up your organization and the infrastructure and services owned and controlled by your organization. The following features in

Microsoft 365 are on by default, or can be configured, to help mitigate the risk and recover from a successful compromise of the assets in this domain.

Exchange Online

- With single item recovery and mailbox retention, customers can recover items in a mailbox upon inadvertent or malicious premature deletion. Customers can roll back mail messages deleted within 14 days by default, configurable up to 30 days.
- Additional customer configurations of these retention policies within the Exchange Online service allow for:
 - configurable retention to be applied (1 year/10 year+)
 - copy on write protection to be applied
 - the ability for the retention policy to be locked such that immutability can be achieved
- Exchange Online Protection scans incoming email and attachments in real-time both entering and exiting the system. This is enabled by default and has filtering customizations available. Messages containing ransomware or other known or suspected malware are deleted. You can configure admins to receive notifications when this occurs.

SharePoint Online and OneDrive for Business Protection

SharePoint Online and OneDrive for Business Protection have built in features that help protect against ransomware attacks.

Versioning: As versioning retains a minimum of 500 versions of a file by default and can be configured to retain more, if the ransomware edits and encrypts a file, a previous version of the file can be recovered.

Recycle bin: If the ransomware creates a new encrypted copy of the file, and deletes the old file, customers have 93 days to restore it from the recycle bin.

Preservation Hold library: Files stored in SharePoint or OneDrive sites can be retained by applying retention settings. When a document with versions is subject to retention settings, versions get copied to the Preservation Hold library and exist as a separate item. If a user suspects their files have been compromised, they can investigate file changes by reviewing the retained copy. File Restore can then be used to recover files within the last 30 days.

Teams

Teams chats are stored within Exchange Online user mailboxes and files are stored in either SharePoint Online or OneDrive for Business. Microsoft Teams data is protected by the controls and recovery mechanisms available in these services.

Domain 2: Service level controls

The second domain is the people that make up Microsoft the organization, and the corporate infrastructure owned and controlled by Microsoft to execute the organizational functions of a business.

Microsoft's approach to securing its corporate estate is Zero Trust, implemented using our own products and services with defenses across our digital estate. You can find more details about the principles of Zero Trust here: Zero Trust Architecture.

Additional features in Microsoft 365 extend the risk mitigations available in domain 1 to further protect the assets in this domain.

SharePoint Online and OneDrive for Business Protection

Versioning: If ransomware encrypted a file in place, as an edit, the file can be recovered up to the initial file creation date using version history capabilities managed by Microsoft.

Recycle bin: If the ransomware created a new encrypted copy of the file, and deleted the old file, customers have 93 days to restore it from the recycle bin. After 93 days, there's a 14-day window where Microsoft can still recover the data. After this window, the data is permanently deleted.

Teams

The risk mitigations for Teams outlined in Domain 1 also apply to Domain 2.

Domain 3: Developers & service infrastructure

The third domain is the people who develop and operate the Microsoft 365 service, the code, and infrastructure that delivers the service, and the storage and processing of your data.

Microsoft investments that secure the Microsoft 365 platform and mitigate the risks in this domain focus on these areas:

- Continuous assessment and validation of the security posture of the service

- Building tools and architecture that protect the service from compromise
- Building the capability to detect and respond to threats if an attack does occur

Continuous assessment and validation of the security posture

- Microsoft mitigates the risks associated with the people who develop and operate the Microsoft 365 service using the principle of **least privilege**. This means access and permissions to resources are limited to only what is necessary to perform a needed task.
 - A Just-In-Time (JIT), Just-Enough-Access (JEA) model is used to provide Microsoft engineers with temporary privileges.
 - Engineers must submit a request for a specific task to acquire elevated privileges.
 - Requests are managed through Lockbox, which uses Azure role-based access control (RBAC) to limit the types of JIT elevation requests engineers can make.
- In addition to the above, all Microsoft candidates are pre-screened prior to beginning employment at Microsoft. Employees who maintain Microsoft online services in the United States must undergo a Microsoft Cloud Background Check as a prerequisite for access to online services systems.
- All Microsoft employees are required to complete basic security awareness training along with Standards of Business Conduct training.

Tools and architecture that protect the service

- Microsoft's Security Development Lifecycle (SDL) focuses on developing secure software to improve application security and reduce vulnerabilities. For more information, see Security and Security development and operations overview.
- Microsoft 365 restricts communication between different parts of the service infrastructure to only what is necessary to operate.
- Network traffic is secured using extra network firewalls at boundary points to help detect, prevent, and mitigate network attacks.
- Microsoft 365 services are architected to operate without engineers requiring access to customer data, unless explicitly requested and approved by the customer. For more information, see How does Microsoft collect and process customer data.

Detection and response capabilities

- Microsoft 365 engages in continuous security monitoring of its systems to detect and respond to threats to Microsoft 365 Services.

- Centralized logging collects and analyzes log events for activities that might indicate a security incident. Log data is analyzed as it gets uploaded to our alerting system and produces alerts in near real time.
- Cloud-based tools allow us to respond rapidly to detected threats. These tools enable remediation using automatically triggered actions.
- When automatic remediation isn't possible, alerts are sent to the appropriate on-call engineers, who are equipped with a set of tools that enable them to act in real time to mitigate detected threats.

Recover from a ransomware attack

For the steps to recover from a ransomware attack in Microsoft 365, see Recover from a ransomware attack in Microsoft 365.

Additional ransomware resources

Key information from Microsoft

- The growing threat of ransomware, Microsoft On the Issues blog post on July 20, 2021
- Human-operated ransomware
- Rapidly protect against ransomware and extortion
- The latest Microsoft Security Intelligence Report (see pages 22-24)
- **Ransomware: A pervasive and ongoing threat** report in the **Threat analytics** node of the Microsoft Defender portal (see these licensing requirements)

Microsoft 365

- Deploy ransomware protection for your Microsoft 365 tenant
- Recover from a ransomware attack
- Protect your Windows 10 PC from ransomware
- Handling ransomware in SharePoint Online

Microsoft Defender XDR

- Find ransomware with advanced hunting

Microsoft Azure

- Azure Defenses for Ransomware Attack
- Backup and restore plan to protect against ransomware
- Help protect from ransomware with Microsoft Azure Backup (26-minute video)
- Recovering from systemic identity compromise
- Advanced multistage attack detection in Microsoft Sentinel
- Fusion Detection for Ransomware in Microsoft Sentinel
- Ransomware protection in Azure
- Prepare for a ransomware attack
- Detect and respond to ransomware attacks
- Azure features & resources that help you protect, detect, and respond
- Azure backup and restore plan to protect against ransomware

Microsoft Defender for Cloud Apps

- Create anomaly detection policies in Defender for Cloud Apps

Microsoft Security team blog posts

- 3 steps to prevent and recover from ransomware (September 2021)
- Becoming resilient by understanding cybersecurity risks: Part 4—navigating current threats (May 2021)

 See the **Ransomware** section.

- Human-operated ransomware attacks: A preventable disaster (March 2020)

 Includes attack chain analyses of actual attacks.

- Ransomware response—to pay or not to pay? (December 2019)
- Norsk Hydro responds to ransomware attack with transparency (December 2019)

Built-in virus protection in SharePoint Online, OneDrive, and Microsoft Teams

Article • 06/09/2023 •

Applies to: ✅ Exchange Online Protection, ✅ Microsoft Defender for Office 365 plan 1 and plan 2

> **Tip**
>
> *Did you know you can try the features in Microsoft 365 Defender for Office 365 Plan 2 for free?* Use the 90-day Defender for Office 365 trial at the **Microsoft 365 Defender portal trials hub**. Learn about who can sign up and trial terms **here**.

Microsoft 365 uses a common virus detection engine for scanning files that users upload to SharePoint Online, OneDrive, and Microsoft Teams. This protection is included with all subscriptions that include SharePoint Online, OneDrive, and Microsoft Teams.

> **Important**
>
> The built-in anti-virus capabilities are a way to help contain viruses. They aren't intended as a single point of defense against malware for your environment. We encourage all customers to investigate and implement anti-malware protection at various layers and apply best practices for securing their enterprise infrastructure.

What happens if an infected file is uploaded to SharePoint Online?

The Microsoft 365 virus detection engine scans files asynchronously (at some time after upload). If a file has not yet been scanned by the asynchronous virus detection process, and a user tries to download the file from the browser or from Teams, a scan on download is triggered by SharePoint before the download is allowed. **All file types are not automatically scanned**. Heuristics determine the files to scan. When a file is found to contain a virus, the file is flagged.

Here's what happens:

1. A user uploads a file to SharePoint Online.
2. SharePoint Online, as part of its virus scanning processes, later determines if the file meets the criteria for a scan.
3. If the file meets the criteria for a scan, the virus detection engine scans the file.

4. If a virus is found within the scanned file, the virus engine sets a property on the file that indicates the file is infected.

What happens when a user tries to download an infected file by using the browser?

By default, users can download infected files from SharePoint Online. Here's what happens:

1. In a web browser, a user tries to download a file from SharePoint Online that happens to be infected.
2. The user is shown a warning that a virus has been detected in the file. The user is given the option to proceed with the download and attempt to clean it using anti-virus software on their device.

To change this behavior so users can't download infected files, even from the anti-virus warning window, admins can use the *DisallowInfectedFileDownload* parameter on the **Set-SPOTenant** cmdlet in SharePoint Online PowerShell. The value $true for the *DisallowInfectedFileDownload* parameter completely blocks access to detected/blocked files for users.

For instructions, see Use SharePoint Online PowerShell to prevent users from downloading malicious files.

Can admins bypass *DisallowInfectedFileDownload* and extract infected files?

SharePoint admins and global admins are allowed to do forensic file extractions of malware-infected files in SharePoint Online PowerShell with the Get-SPOMalwareFileContent cmdlet. Admins don't need access to the site that hosts the infected content. As long as the file has been marked as malware, admins can use **Get-SPOMalwareFileContent** to extract the file.

For more information about the infected file, admins can use the **Get-SPOMalwareFile** cmdlet to see the type of malware that was detected and the status of the infection.

What happens when the OneDrive sync client tries to sync an infected file?

When a malicious file is uploaded to OneDrive, it will be synced to the local machine before it's marked as malware. After it's marked as malware, the user can't open the synced file anymore from their local machine.

Extended capabilities with Microsoft Defender for Office 365

Microsoft 365 organizations that have Microsoft Defender for Office 365 included in their subscription or purchased as an add-on can enable Safe Attachments for SharePoint, OneDrive, and Microsoft Teams for enhanced reporting and protection. For more information, see Safe Attachments for SharePoint, OneDrive, and Microsoft Teams.

Related articles

Malware and ransomware protection in Microsoft 365

Turn on Safe Attachments for SharePoint, OneDrive, and Microsoft Teams.

Microsoft 365 vulnerability scanning and remediation

Article • 04/04/2023

An organization's security posture isn't achieved through a one-time effort, it requires consistent evaluation and maintenance. Systems deemed secure at the time of deployment are soon considered insecure due to newly identified vulnerabilities in the ever-evolving threat landscape. For this reason, systems and baselines should be continuously evaluated and updated through standardized processes enforced across the enterprise. Microsoft 365 routinely scans all its systems and applicable network devices for known vulnerabilities and centrally tracks their remediation in a timeframe directly related to their severity.

Scanning tools and approach

Microsoft 365 uses a third-party scanning tool utilizing two different techniques to ensure full coverage of its assets: host-based and network-based. For all servers and network accessible devices capable of running a host agent, host-based scanning tools are used, and scan results are uploaded to a central location for reporting. All assets that can't run the local agent, such as network devices and a small subset of bare metal servers, are captured by network scanning tools.

For host-based scanning, agents are deployed in a holistic manner as part of the Microsoft 365 service onboarding process and authentication rates are calculated daily to ensure asset scanning coverage with an authentication rate target of 95%. A snapshot of the host containing metadata is created and stored within a sandbox environment where it's then assessed for vulnerabilities. By running the scan in a sandbox environment instead of on host machines, we mitigate the risk that an attacker can perform a supply chain attack.

Scans evaluate hosts and network devices for missing patches, configuration issues, and application vulnerabilities using up-to-date vulnerability information from industry sources such as the Microsoft Security Response Center (MSRC), the NIST National Vulnerability Database, and the MITRE Common Vulnerabilities and Disclosures Database. When new patches and vulnerabilities are discovered by the scanning process, they're added to Microsoft's internal vulnerability database. In addition to vulnerabilities gathered during scanning, the vulnerability database contains data from the MSRC, vendor web sites, and third-party vulnerability detection services.

Host-based scans are performed daily, network scans are performed weekly, and each is configured to search for all known vulnerabilities in the signature database. Scanning processes include a check to ensure that updates have occurred before each use and results are sent to Azure Data Explorer for central storage. From there, scan results are aggregated, prioritized for triage, and reported on via internally developed dashboards.

Reporting

Microsoft 365 Security uses automated reporting to compare scan results over time, displaying new vulnerabilities when they're found. These reports are updated daily and are available to authorized personnel via the Patching and Vulnerability Compliance (PAVC) dashboard. The PAVC dashboard is the source of truth for tracking and reporting all Microsoft 365 vulnerability data. If requested, service teams have the ability to create additional dashboards using PAVC data if desired.

The dashboard displays:

- Live vulnerability metrics including host counts.
- Recommended solutions for service teams
- Metrics for scan quality and compliance with remediation requirements
- A list of all active vulnerabilities.

Remediation

Security bugs are created for each vulnerability reported on in the PAVC dashboard to ensure they're remediated effectively across Microsoft 365. These bugs are then triaged to engineering teams who own affected system. In cases of false positive detections, The Microsoft 365 PAVC team is responsible for working with third-party scanning vendors to ensure resolution. Resolutions may involve modifying detections or creating internal solutions to false positives.

Vulnerability remediation actions vary and include making updates to base images, applying patches, and deploying configuration changes to production servers following change management procedures. Depending on the severity of the vulnerability, patching is required within 30 days for high-risk vulnerabilities, 90 days for moderate vulnerabilities, and 180 days for low-risk vulnerabilities.

Microsoft 365 engineering teams can also file exceptions for vulnerabilities that can't or shouldn't be remediated. The Microsoft 365 PAVC Exception Review Board then reviews these requests and either approves or rejects them based on security and compliance

needs. If rejected, service teams continue with the remediation process. If accepted, the exception is processed so that it isn't flagged for that system during future scans.

Asset coverage

To ensure remediation across all systems in Microsoft 365, PAVC compares its results to a complete inventory of all physical and virtual assets.

A list of virtual assets is maintained in an internal Azure tool which is automatically updated when new assets are created and deployed. Physical Assets are maintained by the individual Service Teams that own them.

PAVC uses automated scripting and queries to centralize inventory data into PAVC tooling and conducts a monthly review to ensure scripts continue to capture a complete and accurate list. This allows the PAVC team to track patching coverage and ensure that vulnerability remediations are comprehensive.